

# LUNA Lab complete user guide

Dongduk Women's University · Data Science · LUNA Lab

[lunablab-ai.com/workspace/help/](https://lunablab-ai.com/workspace/help/)

Reviewed 2026-10-12

## Start with your task

### Research with AI in a space

Research space guide → Research goals, literature and strategy → Shared literature, research code and experiments → Connect ChatGPT and execute actions

### Manage your research day

Today's research: planning and daily review → My work: unified browsing and management → Actions and follow-up: recommend, compose, approve and review → Report requests, responses and templates

### New researchers

Getting started → Membership, invitations and managers → Profile, language and translation → Materials, classification and references

### Collaborate on research

Research messenger and AI consultation → Grant wiki · organize important knowledge → Calendar and Google event creation → Share and cite links

### Use AI and reports

Generate, schedule and review reports → Personal AI memory and source coverage → Create and download AI files → Connect ChatGPT and execute actions

### Lab and project managers

Collection and synchronization → Manage material AI analysis → Drive preservation and research context → AI assessment criteria, evidence window and refresh

# Contents

## Teaching and grading

1. Grading and assessment

## Getting started

2. Getting started
3. Membership, invitations and managers
4. Profile, language and translation

## Research method administration

5. Research skill management

## Daily research and work

6. Today's research: planning and daily review
7. My work: unified browsing and management
8. Actions and follow-up: recommend, compose, approve and review
9. Report requests, responses and templates
10. Research goals, literature and strategy
11. Shared literature, research code and experiments
12. Research workflow, evidence and reproducibility readiness
13. Research space guide

## Grants, projects and connections

14. Operating grants and projects
15. Grant wiki · organize important knowledge
16. Collection and synchronization
17. Connect a GitHub repository
18. Connect GitHub Projects and manage work
19. Connect Notion

20. Connect OneNote research notes

21. Connect Google Drive

22. Connect Zotero references

## Materials and calendar

23. Materials, classification and references

24. Upload materials and choose cloud files

25. Change requests and trash

26. Manage material AI analysis

27. Calendar and Google event creation

## Reports and publishing

28. Generate, schedule and review reports

29. Reproducibility results, Python and reports

30. Write reports, control privacy and receive email

31. Reusable report instructions

32. Publish reports, materials and events to Notion

33. Print, PDF and document exports

## Conversations, AI and collaboration

34. Research messenger and AI consultation

35. Compose, cite and delete messages

36. Personal AI memory and source coverage

37. Create and download AI files

38. Connect ChatGPT and execute actions

39. Share and cite links

40. Discussions, follows and publications

41. Favorites, trash and activity review

## Researchers and lab news

- 42. View related GitHub work on a profile
- 43. AI assessment criteria, evidence window and refresh
- 44. Personal Zotero publications and the lab bibliography
- 45. Read and publish the AI activity magazine
- 46. Write announcements and request AI drafts

## Timebank, notifications and support

- 47. Timebank: deposits, transfers and GPU usage
- 48. Profile and notifications
- 49. Drive preservation and research context
- 50. Troubleshooting and support

# 1. Grading and assessment

Account-scoped grading, inline answers and rubrics, instructor review, XLS and private PDFs.

[Open this guide online](#)

## Graders and instructors

Open Grading in the profile menu or /workspace/grading/ with your own LUNA account. The instructor assigns an exact login ID, questions and students. Read answers, solutions and rubrics alongside scoring. Public feedback and internal notes remain separate; instructor review, corrections and final versions retain their history.

## Files, security and costs

An approved final snapshot produces a SmartClass XLS and individual PDFs. File creation does not prove LMS delivery; verify the saved LMS result. Reports require an explicit account mapping and instructor release. Never distribute a shared attachment containing all reports. Every request checks current authority; CSRF, revisions, replay protection and finalization locks remain enforced. Records follow encrypted off-server database backups, recoverable through the last verified checkpoint. Grading invokes no AI and incurs no AI usage charge.

## Fast question navigation

Use Save and next question or Ctrl+Enter (Command+Enter on Mac). After a successful save, focus moves to the first criterion of the next assigned question; the last question continues to the next student in the current order. Open the question number to see this student's assigned entry/review-request statuses and jump directly. Same question · next student retains the question. Missing input, failed saves and conflicts retain the current input and stop navigation. Entered status does not mean instructor review or final approval.

## Review, backup and delivery

Sort by name, student ID, score or request count; filter to students and questions needing instructor review. Download individual final PDFs or a complete PDF/XLS ZIP. Full instructor backups also contain keys, internal notes, assignments and history and must never be shared with students. Encrypted Drive checkpoints and course-specific PC sync preserve data through the last verified timestamp. Codex delivery verifies the latest final snapshot, course and individual recipient, then checks the saved SmartClass result; authentication and uncertain outcomes pause delivery.

## 2. Getting started

Start with membership, your profile and your first research material.

Open this guide online

### 1. Prepare your accounts

Prepare your university email, platform username and relevant GitHub, Notion and Google accounts. Emails need not match across services, but each account must be able to open its sources. Never put passwords or tokens in chat or reports. Choose dedicated project folders and pages, separate from personal content.

### 2. Register and join

1. Register with the required account details and consents. If you followed an invitation, accept it after registration or sign-in.
2. Otherwise, find the grant or project under Grants → Request participation. Check Notifications for the decision.
3. Registration, lab membership, grant participation and child-project participation are separate. You need active access to the relevant scope.
4. After approval, check filters and collection status if materials are missing. A matching author name or email does not grant access.

### 3. Profile versus source connections

Use View profile → Edit my profile for your name and biography and Change photo for a JPG, PNG or WebP image. A GitHub username helps identify authors; it is not repository authentication. Open your approved project under Research grants, then External source settings. If management controls are absent, ask the project lead to configure a connection.

### 4. Connect a small sample first

Start with one report or a dedicated project folder. Enable the connection, save and click Collect now once. Check collection time, count and status, then compare the collected body with the original under Materials. Once verified, enable periodic collection, for example every 60 minutes. Periodic collection is separate from scheduled AI report generation.

### 5. First-week checklist

Monday: verify membership, profile and source connections. During research: record date, author, objective, work, results, blockers and next steps in the original tool. Before meetings: check the event and material links. Afterwards: record decisions, owners and deadlines. At week-end: check collection failures or gaps and review evidence-based reports with your lead.

## Daily menu and suggested reading order

1. The top menu is Overview → Reports → Grants → Calendar → Materials → Researchers → Timebank → Notifications.
2. New researchers can read Membership, Profile/notifications, Projects, Materials, Related GitHub work and Timebank in that order.
3. Administrators should prepare membership and source connections before configuring reports, schedules, instruction templates, AI assessment, magazine and Timebank.
4. Expand a topic group in the contents, choose a chapter and use its section links. Search for a task such as deposit, Only me, deadline or tags.
5. Download the full guide in Korean or English; it uses the same content as the website. Missing buttons may reflect your role or disabled features.

## Stay signed in and use Overview

1. Select Keep me signed in when logging in on a device you control. Logout, password changes and revoked access still end access.
2. On Overview, start with the magazine, popular pages and recent reports. Popularity is based on activity on pages you may read.
3. Announcements and network news show previews; expand for more. The Following filter collects updates related to researchers you follow.
4. Some materials, profile sections and discussions load independently. Wait or use the offered retry/full-content link instead of repeatedly reloading the entire page.

## Mobile profile and chat icons

On mobile, open profile, settings and sign out from the profile icon to the right of notifications and help. The header omits your username. The left menu holds primary navigation and language selection. The floating speech-bubble icon opens conversations.

## Choose your start page

1. Choose Overview, Today's research, My work, Conversations, AI consultation or a specific grant/project in Start page settings on Settings or Edit profile. It applies only to your account.
2. Choose a conversation list/accessible channel, or search for the actual grant/project. Saving does not navigate away; use Open now to try it.
3. The logo and default platform entrance use your preference. The Overview menu, direct material/notification links and an explicit login destination still open their intended screen.

4. If the target is removed or no longer accessible, the app falls back to Overview. A start preference does not grant access.
5. Opening AI consultation as your start page does not submit an AI request or use allowance. Choose Overview to return to the default.

## Read help alongside your work

1. Select ? Help in the header to open relevant guides and search. A guide link beside a feature opens that explanation directly. Opening or searching help keeps your current form values and selections.
2. On wide screens, use your working page and the right help panel together. On small screens, help fills most of the screen; Close or Esc returns to your work. Help and conversations open one at a time without clearing a message being composed on the same page.
3. Search in Korean or English and read results in the current interface language. Moving to another working page automatically shows its guidance. Pin guide keeps the guide you are reading; Help for this page or the header Help button returns to the new screen guidance. Previous guide and All topics remain available.
4. Navigation between working pages follows existing loading and saving behavior. Save working edits before leaving a page. Help does not automatically save your working form.
5. Open in a new tab or Full help center provides a larger reading view. If loading fails, retry while keeping current inputs. Search covers public guides only; it does not search or send private research material or form contents to AI.

## Ask AI about the current page

1. In the right help panel, choose Ask AI. Check the current screen and AI lab in Settings, then type a question in the bottom composer. On mobile, the conversation takes most of the screen; lab selection, usage and Clear conversation are under Settings. The composer grows to three lines. While reading earlier messages, use New answer to reach the latest response. Replies follow the Korean or English interface language.
2. Only your question, recent help conversation, screen type and public handbook are sent to AI. Page research content, files and form drafts are not sent automatically. Do not paste secrets or research originals into your question.
3. AI permission and an enabled lab configuration are required. Only your personal AI allowance applies; shared lab analysis attempts are not used. You can stop waiting; an attempt already started may still count toward usage.
4. Open the cited guides in the same panel. AI does not inspect live server state or save or modify work. Conversations are private to your account, retained for seven days, and removable with Clear conversation.

5. When no answer is available, or after This did not resolve my issue, choose Review a Customer support draft. In the new tab, review the question, checked guides and AI response, add reproduction steps and expected and actual results, then choose Submit request yourself. Opening a draft does not submit a ticket.

## Open mobile navigation

1. On small screens, the three-line button at the top left opens a navigation panel. Notifications and Help remain directly accessible in the header.
2. Choose a workspace page and language in the panel. Open Profile, Settings and Announcements from the header profile icon to the right of notifications and help.
3. Close, click outside or press Esc to return. Opening and closing the menu preserves form inputs and scroll position. Save edits before navigating to a different page. Navigation, Help and Conversations open one at a time.

## 3. Membership, invitations and managers

Request access, invite colleagues and manage membership.

Open this guide online

### 1. Request membership

1. Sign in → Grants → Join.
2. Find a grant or project and submit a request. Grant and child-project memberships are separate.
3. Check Notifications for approval. After joining a grant, you can also request a child project from the grant homepage.
4. Reopen Grants → your project after approval. If unavailable, ask the manager to check both approval and active participation.

### 2. Invite a colleague as a manager

1. Open the grant or project → Participants and invitations.
2. For an existing member, choose the researcher under Add participant, review the role and management permission, then save.
3. For a new member, generate an invitation link, copy it and send it yourself. The recipient follows the link, signs up or signs in, then accepts.
4. Registration alone is not acceptance. Confirm the recipient in the participant list. Issue a new link if expired. The platform does not automatically email invitations.

### 3. Approve, change roles or remove participation

1. Open a join-request notification, check its scope and applicant, then approve or reject.
2. Use the participant list to assign or remove management permission. Grant managers manage that grant and its child projects; project managers manage their project.
3. Removing participation ends access to that scope; it does not delete the account or remove unrelated memberships.
4. If controls are missing, check your management role. Lab management is available to platform administrators.

## Pending management requests

1. Platform administrators can open the profile menu to see the awaiting-reply count next to Support. It counts unresolved new tickets and tickets whose latest reply came from the requester. An administrator reply clears that waiting state; a requester follow-up adds it again.
2. The Review change requests card in Lab management counts pending metadata and deletion requests you currently have permission to review. Decided requests and inaccessible targets are excluded.
3. Each Membership and invitations card counts pending membership approvals for that lab. Project/grant applications and unused invitation links are not added.
4. Zero counts are hidden. Reopen the profile menu or management page to load current counts; returning to the page also refreshes them. These counts reflect pending work, separately from unread notifications.

## Permanently delete a member account

1. As a platform administrator, open Lab management → Members and invitations. Choose Permanently delete account on the member card. Suspension is reversible; permanent deletion is not.
2. Review the username, email, all lab memberships, connections/calendar credentials/report schedules to stop, private AI conversation count and Lumen balance. This deletes the platform account across all labs, rather than removing only one lab membership.
3. You cannot delete your own account here. Transfer the last active PI role first. Transfer owned grants/projects through Researchers → the successor's management page → scope → transfer responsibility. Settle any Timebank balance before deletion.
4. Enter the target username exactly and your own administrator password. Check the irreversible-deletion acknowledgement, then submit. Review again if ten minutes elapse or the target changes. Cancel performs no deletion.
5. The original login row, memberships, profile/personal settings, private AI conversation text and memory, favorites, follows, notifications and personal Zotero credentials are removed. Existing sessions lose access. Researcher references become inactive anonymous identities.
6. Research reports, shared documents, comments and ledger history retain an anonymous, non-login author reference. Names and source text within research bodies/attachments are not rewritten. Connections created by the account lose credentials and stop collecting; report schedules stop. Another manager must reauthorize needed shared connections.
7. This affects the current platform database. Previous operational backups and original accounts/material on GitHub, Google, Notion and other services are not deleted by this button. Review the complete impact before submitting.

## Check join and change-request notifications

1. Open join or metadata-change requests from Notifications and check the scope and applicant.
2. Review and approve or reject with the required current authority. An administrator label alone does not replace a scope-specific authorization check.
3. If no alert appears, check relevant notification categories, administrator policy and approval rights. Reading or deleting a notification does not itself decide the request.

## 4. Profile, language and translation

Choose a language and compare on-device translation with AI copies.

Open this guide online

### 1. Profile and photo

1. Open the profile icon → View profile → Edit my profile.
2. Enter relevant name, email, phone, affiliation, job title and biography details, save and check your profile homepage.
3. In Edit my profile, use Profile photo to upload JPG, PNG or WebP and refresh to verify both profile and navigation images.
4. Changing your username can affect your profile URL. A GitHub username is not repository authentication or an access grant.

### 2. Choose Korean or English

1. Choose 한국어 or English in the header. Menus and help change language; internal navigation and searches preserve it.
2. Copy the URL including lang=KO or lang=EN to share the selected language. Recipients still need access.
3. Display translation does not rewrite the stored original. Use Show original and verify numbers, formulas, names and citations.
4. On-device translation and AI copies have different requirements and allowances. Choose the appropriate workflow below.

### 3. English lists and low-cost translation

1. Home, Materials, Reports and Calendar check saved translations of titles and short excerpts first. Prepared English editions remain authoritative; originals are never overwritten.
2. Supported desktop Chrome uses free on-device translation first. Choose Enable free translation for initial setup. Results are saved privately to your account and reused for free on mobile. If the on-device API is unavailable on mobile, Chrome's own Translate menu is another free option.
3. Choose Translate missing text with AI and review the scope and personal allowance. One execution covers at most 20 visible text segments / 6,000 characters. Failed actual executions also count; cache reads and duplicate waiting requests do not. The shared lab analysis allowance is not used.

4. Settings → List translation settings offers Free only / Ask each time (default) / Automatic within an explicitly approved daily limit. The separate default limit of 3 executions and your personal allowance both apply, resetting at midnight in Seoul. There is no automatic upgrade to a more expensive model.

5. AI copies can be reused by readers with access to the same source; only the first requester uses an allowance. Browser copies are private to you. Changed source text needs a fresh copy. Compare with View original and check numbers and names. Forms, links, attachments and original records stay unchanged.

## Links that choose the interface language

Add ?lang=en to a URL to open it in English, or ?lang=ko for Korean. Example:

<https://lunabank-ai.com/workspace/bank/usage/?lang=en>. Uppercase EN and KO also work. If the URL already has query parameters, append &lang=en. The selection applies before login and remains after login and on subsequent pages. The header English button adds lang=EN to the address so you can copy the English link. The marker stays in the address when navigating or searching. The Korean button or a ?lang=KO link selects Korean and keeps lang=KO on subsequent pages. Source content and access permissions are unchanged.

## 5. Translate material and report details

1. Choose Read in English below a material or report title. Supported desktop Chrome translates its displayed title, body, tables, notes and analysis on-device.

2. Use Show original to restore it, or Cancel while processing. Code, inputs, original attachments and downloads stay unchanged. No LUNA AI allowance is used.

3. Translation options → AI-translated copy is a separate request. Review its language and transmission confirmation; current AI permission and daily allowance apply.

4. If unsupported or unsuccessful, read the original or use the offered browser guidance. Verify research conclusions against the source.

## Automatic translation in daily work and requests

1. In English mode, supported reading areas in daily notes, My work lists/details and report requests/responses use the existing automatic translation flow. Compare View original when needed.

2. Free on-device translation and cached copies depend on browser support, language-model availability and saved copies. Original text can remain visible while unavailable or preparing.

3. Editors contain original text. Viewing a translation does not rewrite saved notes, GitHub content or another researcher's writing.

4. Links, code, mentions and numbered references keep their destinations. Inspect sources and the actual outgoing preview before sharing translated content.

5. Display language and AI generation instructions are separate. Check generation instructions for desired AI wording. Browser translation uses no LUNA AI allowance.

## 5. Research skill management

Administrator research guidance, candidate evaluation, promotion, rollback and backup.

[Open this guide online](#)

### 1. Applied guidance

Research-kind baselines combine with approved domain, topic, space and work guidance. Execution preparation freezes exact versions; promotion does not rewrite ongoing results. Explicit revocation and current source permissions are still checked. Skills are not model-weight training or proof of research success.

### 2. From candidate to approval

Only current platform administrators manage candidates, contract evaluation, promotion, rollback and revocation. Preview each exact change before applying it. Method improvements actually delivered in an answer may be preserved only as unapproved candidates in that private work. Missing proposals create no candidates, and saving invokes no additional AI. AI never replaces approved guidance automatically.

### 3. Evaluation and cost

Contract evaluation compares the active and candidate instructions against fixed cases and records results, differences and hashes. It does not measure model-answer quality or experimental success. Before promotion, administrators review sources, scope, limitations, cost and their review rationale. Requesting improvements within a paid research answer adds tokens to its existing usage. Separate paid quality evaluation is not started automatically.

### 4. Privacy and recovery

An administrator role does not grant access to another person's private conversations or memories. Private candidates remain bound to their owner and current source permissions and are never automatically copied into global guidance. Write a separate generalized candidate and explicitly review it for private information and secrets. Authorized version exports do not replace complete encrypted recovery backups. Instructions sent to external ChatGPT cannot be guaranteed hidden.

### Keep only the know-how you choose from a discussion

In the same work composer, ask: "Separate the literature-analysis and manuscript-writing know-how from this discussion. Propose only the changes to each existing method." Only delivered, saved requests and results are evidence. External ChatGPT or Codex can prepare structured proposals using returned evidence IDs and current permissions, then use the existing reviewed save contract. On the proposal page, select only the drafts you want to keep and review the save. Unselected proposals and active versions remain unchanged. Preparing or saving a draft does not validate quality or approve v2.

## Apply several methods together

Use the applied know-how count below the work title, or work menu → Know-how settings, to manage defaults. No empty badge appears for zero items. Search names, handles or activities and review exact versions and roles before saving. Add to a work on the independent detail page lets you choose a space and work and pin an exact version. Type @ in a follow-up, new research request or personal AI consultation to search researchers and know-how together. Choosing know-how adds it only to that request without changing defaults. [[ source citations keep their existing behavior. Multiple selections never auto-merge assets or adopt newer versions. Web, consultation and ChatGPT handoff check the same IDs and current permissions.

## Split, merge, originals and shared evidence

One discussion can produce independent proposals that reference the same segment as shared evidence. Drafts start privately. Separate paper-specific corrections and unverified scientific conclusions from reusable guidance. Change organization prepares a split or merge request in the existing composer for review and sending. A merge proposes a new asset linked to exact parent versions. It does not delete or overwrite the versions, history or active state of assets A and B. Changing how a proposal is organized requires a newly reviewed proposal, not an unreviewed browser-side rewrite.

## Version inheritance and evaluation

Derived work inherits the exact selected versions of independent assets. For existing work-specific methods, the collapsed Work know-how controls retain version and candidate selection in the original work. Review generalized instructions and create a separate asset before reuse in other work. Execution applies modules relevant to the current stage, such as literature, synthesis, writing or review, while retaining links to other stages. New versions are not adopted automatically. Initial approval reviews the actual instructions and their generalization; it does not claim measured quality improvement. Improvement candidates require actual runs with matched questions, inputs, models and effort, held-out cases and human review before adoption. A candidate trial is a one-off selection that does not change the default and uses the normal AI allowance.

## Protection, storage and execution boundaries

An independent detail page shows introduction, ID, handle, authors, usage, current permissions and versions. Editing the introduction or image link does not change execution guidance; image access is checked separately. The default direction is protected guidance. Write a draft or propose a direction improvement, review it and select which drafts to keep. Saving an introduction or candidate does not change the active version. Ordinary catalogs and proposal cards exclude protected instructions. Source review requires current permission and AI-provider processing permission is separate. Reads, selections and structured proposal saves do not run AI, while sending analysis or reorganization requests may use your allowance. Independent external analysis does not prove protected-service execution. Instructions belong in encrypted backups, not code or static files. AI quality, actual ChatGPT tool availability and server speed require separate validation.

## 6. Today's research: planning and daily review

Plan, edit, draft with AI, act and review your day.

Open this guide online

### 1. Start your morning

1. Open the profile menu → Today's research. Check the date, then review assignments, deadlines and events in Today's direction. Use the date picker or arrows to view other daily records.
2. Choose Write or Edit beside Today's goal, enter your priorities and save. AI plan draft is optional; writing your own plan is fully supported.
3. Expand Select focus tasks to choose up to three. Use Act on a plan item to prepare a comment, message or other next step. The sections below explain applying drafts and executing actions.
4. Check new requests and follow-ups in Reports & actions. A red number counts new requests, not unfinished assignments or executions.
5. View all work opens My work. A written plan and GitHub task state are separate; check the original task state when work is finished.

### 2. Understand dates, counts and evidence

1. Selecting a date opens that day's personal notes. It does not reconstruct a complete historical snapshot of currently collected work. Check the date and collection time of each source.
2. Unfinished assignments, deadlines, events and evening checks measure different things. A 3/5 check-in means three checklist boxes are checked, not three completed tasks or submitted reports.
3. Expand schedule/collection status or observed participation to inspect evidence and limits. A scheduled event does not prove attendance; missing observations do not prove inactivity.
4. For missing work, check your GitHub ID, assignment, active grant/project participation and source collection. Inaccessible sources are omitted.
5. Recent records lets you review earlier private plans and outcomes. If the wrong date is selected, choose Today before entering new notes.

### 3. Write with Markdown, mentions and citations

1. Goals, results, blockers and next steps initially appear in reading mode. Choose Edit beside a field to change its original text.
2. Expand editing tools for headings, lists, bold text and links, then check the preview. Short sections with headings make both reading and action selection easier.

3. Type @ and select an exact researcher username. Type [[ to search and select an accessible document. Typing a similar name alone is not the same as selecting a resolved reference.
4. Save updates your private note. Mentions and citations in a private note do not send messages or grant source access.
5. Save and act first saves your changes, then opens section selection. If another tab has changed the record, resolve the conflict before saving again. Add files for outgoing messages in the action editor.

#### 4. Generate morning and evening drafts separately

1. Use AI plan draft in the morning and AI review draft in the evening. They have different purposes and destination fields; review drafts prepare outcomes, blockers and next steps.
2. Use the arrow beside a generation button for instructions or scope that apply only to this request. Save reusable preferences in AI writing preferences.
3. Check preparation, ready or failure status. Review the prepared text and apply the sections you want. Generation alone neither replaces every saved field nor reports to your PI.
4. When evidence is sparse, distinguish observed activity, your own notes and matters still to check. Add relevant facts yourself rather than presenting unperformed work as an outcome.
5. Check source links and dates. If records or access changed during generation, prepare again with current information. Editing saved text is separate from approving an unapplied draft.

#### 5. Personalize the role, instructions and scope

1. Open AI writing preferences and choose role-aware or another available perspective. PI writing emphasizes supervision, review, decisions and support; researchers' results are not claimed as your own work. Platform administration alone does not imply a PI role.
2. Use common instructions for length, tone and priorities, for example: For each project, list at most three decisions or questions for researchers. This does not change your real role or access.
3. Save separate morning and evening preferences for plans and reviews. Use the generation button's options for a one-time change.
4. Choose organization by researcher, grant or project. Multiple selections within one category use OR; selections across categories use AND.
5. Narrow filters can reduce evidence or exclude private notes without a matching scope. Check filters when evidence is sparse. Restore defaults resets writing preferences, not saved daily records.

#### 6. Group items and inspect numbered evidence

1. Structured AI content offers Group view to change how items are displayed. Switching researcher, grant or project grouping does not run AI again.

2. Expand the short researcher/project context to see whose work is involved. An item may have more than one relevant affiliation.
3. Follow [1], [2] and other numbered references to inspect evidence. A link does not grant access; removed or newly restricted sources cannot be reused.
4. Group view is available for structured drafts. Manually written or extensively restructured text may not offer the same grouping, but saved content can still be used to prepare actions.
5. Changing display order is different from changing generation preferences. Save a preference if future drafts should use that organization.

## 7. Continue with Act and its count

1. Choose Act beside an AI item or a saved manual field. For manual text, select the heading, paragraph or list item to act on. Empty fields have no content to execute.
2. Candidates use accessible citations, links and exact mentions. Inspect and change the destination if the match is wrong or several destinations are possible.
3. The number beside Act counts confirmed publications or personal execution records. It excludes drafts, AI requests, page views, retries and edit history. Select it to see linked actions.
4. Later editing does not overwrite earlier execution content or history. Read any source-changed notice and check the text to use for a new action.
5. Read Actions and follow-up for comments, messages and report handoffs. Pressing Act alone does not publish anything.

## 8. Review outcomes and continue tomorrow

1. Before leaving, record results, blockers/support needed and tomorrow's steps. Separate directly verified outcomes, other people's reports and unresolved questions.
2. In an action, open Results and next check to inspect replies and later activity and save your result. A new evening draft uses accessible actions, observations and personal notes within its evidence window.
3. The five checks record your own review of work registration, evidence, status, reporting needs and next steps. Checking them neither completes GitHub work nor sends a report.
4. Manage unresolved actions with a next-check date. Marking your follow-up finished removes it from future follow-up candidates; it is not research completion or PI approval.
5. For tomorrow's plan, review remaining follow-ups and newly collected evidence. Yesterday's saved review is not automatically rewritten when a new comment arrives.

## 9. Reports, requests and sharing with your PI

1. Use Reports & actions for received reports, incoming requests, sent requests and follow-ups. Request a report appears according to permission; ordinary researchers can answer received requests.
2. Personal goals and results are private by default. To share with a PI, choose the grant and the exact text, then review the actual recipient and content.
3. Shared check-ins appear in LUNA's received reports. A platform notification is created; email and push depend on recipient settings and delivery configuration. Shared does not guarantee email receipt.
4. If no participating grant with another PI is available, check the grant's responsible person and participation. The app does not infer a PI from an arbitrary researcher.
5. Submit a response to a report request from that request's response screen. Saving a personal note or sharing a separate PI check-in does not submit that response.

## 10. Reminders, scheduled drafts and AI usage

1. Set morning/evening times, workdays, time zone and delivery channels in Notification settings. Reminders and scheduled drafts require opt-in; opening the page does not enable automatic delivery.
2. Automatic morning/evening preparation creates drafts. It is different from applying all drafts or sending external comments on your behalf.
3. Check today's AI usage. The consultant limit is five per day for ordinary users and fifty for platform administrators; common personal AI limits and permission still apply.
4. Reading, manual saving, grouping and basic destination matching do not use AI allowance. Optional AI recommendations, wording and advice count according to request and execution policy. Failed attempts that started may count.
5. For pending, failed or limit messages, inspect usage and settings. Missing evidence and changed permissions are different problems; read the message and check pending work before repeating generation.

## 11. Use the same record through ChatGPT

1. Check the ChatGPT connection's additional consent for daily-note saves and research actions. Existing read access does not automatically allow writes.
2. For example: Read my manually written goal today and suggest a question for the cited issue. Show a draft; do not send it. Check the date, section, destination and evidence before choosing.
3. Save ChatGPT's wording as a LUNA action draft, then review the actual actor, destination and body before approval. The web and ChatGPT share records, revisions and execution history.

4. Text composed by ChatGPT is different from requesting LUNA AI generation. Reads and ordinary draft saves do not run extra LUNA AI; explicit LUNA AI requests use its allowance.
5. If tools are missing, check connection metadata and consent and start a new conversation. Use LUNA's web action screen when the host does not expose a required tool. Disconnecting does not undo existing publications.

## 12. Troubleshoot content, controls and translation

1. If Act is missing, first save nonempty content. Manual-field controls and structured AI-item controls appear in different places. Refresh the browser to load the current UI.
2. If Group view is missing, check whether the text is manual or retains structured AI items. Do not delete notes or tasks to recover controls.
3. English mode uses translated reading content with an original view. Browser support, language-model availability and cached translations affect readiness. Editors preserve the original text.
4. Quota errors, sparse evidence, source changes and access changes are different. Follow the specific message, add facts or reselect the destination. For uncertain delivery, check the original and action history first.
5. If needed, contact support with the page address, time, button and error text. Do not include passwords, tokens or private research content.

## AI draft time limits and recovery

1. Daily plan/review drafts use the administrator call timeout.
2. Shared daily allowance, the daily research cap, per-request call budget and timeout are separate. Details and recovery shows applied settings and linked attempts.
3. A received but rejected response is encrypted for up to 24 hours and can be revalidated without AI. Personal notes are not automatically changed.
4. Retry proposes one new call; replaying the same approved review does not duplicate execution. Old failed responses that were never retained cannot be recovered.

## Draft recovery submission and progress

1. Retry review and saved-response revalidation first record a receipt, then check evidence and access in the background.
2. Received, evidence checks, execution queue, AI generation and response validation reflect actual states. Accepted requests survive closing or reloading the page.
3. Confirm the review before accepting a new AI attempt. Saved-response revalidation does not call AI; resending the same request is deduplicated.

4. If disconnected, reload to find the accepted request. Only status reads retry automatically, never failed AI generation. ChatGPT uses the same access, review and call-budget rules.

## Research daily limits and the overall allowance

Use Manage limits in Today's research to open Daily usage limits in Platform AI management.

Administrators choose Follow overall AI allowance or Set a separate cap. Administrators inherit by default; the existing regular-user default remains 5. Inheritance is not extra credit: research shares the personal pool with translation, help and other AI. A separate cap of 0 disables research AI.

Manage my overall daily AI allowance opens per-user overall limits. Attempted runs, queued reservations, completed and failed runs and reset times are distinct.

The overall pool resets at Seoul midnight; a separate research cap resets at midnight in the research timezone. Per-request budgets and per-call timeouts remain separate execution settings. Raising limits never clears history or usage, or automatically retries failures.

In ChatGPT or AI consultation, read `research_usage` via `consultant_account_read`. Only administrators can review and commit execution changes to `research_daily_limit/admin_research_daily_limit`; null inherits the overall pool.

## 7. My work: unified browsing and management

Filter by scope, assignee and deadline; sort, inspect and create work.

[Open this guide online](#)

### 1. Get started with My work

1. Open the profile menu → My work, or Today's research → View all work. This combines currently accessible GitHub work across grants and projects.
2. Choose scope, assignee and deadline, then Open, Overdue, Done or All states. The Me selection uses your linked GitHub assignment identity.
3. Open a title or Details for the body and linked scopes. Use Open on GitHub when you need the current original.
4. Expand Collection status for timestamps and connection warnings. This is collected data, not an exhaustive live query of every GitHub task.
5. Return through Today's research to plan your next steps. Filters do not change your notes or original tasks.

### 2. Choose scopes and assignees

1. Use All scopes or choose a specific grant/project. Distinguish the entire grant from its individual projects.
2. Open the assignee picker and search by name or GitHub ID. Me, All assignees and Unassigned are exclusive choices; several named researchers can be selected together.
3. Assignee refers to actual assignment. To find work you authored, commented on or were mentioned in, also inspect Relation to me under Search & filters.
4. Scope and assignee filters apply together. Widen to All scopes and All assignees for an empty result; this never expands your access rights.
5. If an assignment change is missing, check the original and latest collection time. Matching names alone do not link GitHub identities.

### 3. Combine deadlines, states and detailed search

1. Deadline choices include All, This week, Next week, Next 30 days, Next 90 days, Custom range and No deadline. No deadline finds undated tasks.
2. Expand Search & filters for text, relation, priority and other conditions, then Apply search. Keep it collapsed during ordinary browsing.

3. Entering dates activates a custom range. Provide a start or end date and fix validation messages before applying.
4. Open and Overdue are different filters. A future deadline range can hide overdue work; choose All deadlines and the Overdue tab to inspect it.
5. Active search conditions include collapsed advanced filters. Reset filters when results are unexpected, then add conditions one at a time.

#### 4. Sort, group and interpret counts

1. Choose Due date, Newest created or Oldest created. Creation sorting uses the original GitHub creation date, not import time or latest comment time. Undated items appear last.
2. Switch between List and Board without resetting state tabs, scope, assignees, search or sort. Choose All states for the complete workflow. Group by All, Grant, Project or Assignee. This changes display only, without AI calls or state changes.
3. An issue on multiple boards counts once in the total. Joint assignments or multiple project memberships can appear in several groups.
4. List groups contain the current page; use Previous and Next. Board column counts include all matching work; choose Show more beneath a column for additional cards. To do, In progress and Marked done are the main columns. Review, Blocked/on hold and Unspecified appear when populated. Excluded by filter does not mean there are no tasks. Overdue work is highlighted on cards.
5. On mobile select a column from the compact status row. Expand a group when grouping by grant, project or assignee. A card title opens details; Back to all work restores the board context and position. Use Edit work for an authorized connection in details to change state. Dragging cards to change state is not provided.
6. When boards disagree on state or deadline, compare each membership in Details. One board's Done state does not synchronize every board or close the issue.

#### 5. Save your preferred view

1. Apply the scope, assignee, deadline, state, search, sorting and grouping you want, for example All scopes, Me, All deadlines, Open and Newest created.
2. Use  beside the state tabs → Save current view as default. This is your account's My work preference, not another researcher's view.
3. A shared URL can contain explicit filters. Distinguish its query conditions from your saved default.
4. Use Remove saved default in the same menu to clear the preference. This does not delete tasks or GitHub issues.

5. If participation or access changes invalidate a stored scope, choose currently accessible filters and save again.

## 6. Read and change work details

1. Open Details for the body, issue state, deadline and linked grants/projects. Use Open on GitHub if the body is unavailable.

2. Inspect each board's state, deadline, priority and collection time. A conflict warning means memberships manage the same issue differently.

3. Edit work appears for connections you can manage. Check the destination board, change assignees or fields, and verify the result. Missing controls can indicate access or connection limits.

4. English reading mode uses the existing automatic translation with an original view. Unavailable translation or unsupported browsers may leave the original visible.

5. Back to all work returns to the list context. Copying what you read into a personal plan does not change the original task.

## 7. Create research work

1. Choose New research work, then the destination board. Only boards in grants/projects you can manage are offered.

2. Continue to the work form and fill in title, description, assignees and board fields. This creates work through the selected GitHub connection rather than an independent task source in the combined list.

3. Check the repository and board before creation. Resolve permission or required-field messages in the form or connection.

4. Inspect the result and original link. If delivery is uncertain, check GitHub and history before creating the same work again.

5. If no board is available, ask the connection manager to check active connections, collection and Projects permission. The GitHub Projects guide covers board setup.

## 8. Diagnose empty lists and delayed collection

1. Reset filters and try All states, All assignees and All deadlines. If work is still missing, expand Collection status.

2. Check My GitHub ID settings on your profile. Enter a GitHub username, not a login email or organization name.

3. Recent GitHub changes may appear after the next collection. Follow the warning's connection to inspect authorization and collection results. Repository read access and Projects access can differ.

4. Another researcher's work is limited to what you may currently read. All assignees does not expose every private lab record.

5. For persistent omissions, send support the source URL, time, filters and collection warning. Never send authentication tokens.

## 8. Actions and follow-up: recommend, compose, approve and review

Match destinations, compare AI options, approve actions and review results.

Open this guide online

### 1. Prepare an action from a plan

1. Act on Today's research opens a side panel beside your plan; Full page opens the dedicated screen. Items in the Follow-ups list open as full pages. The top Follow-ups link returns to the previous filters, page and scroll position. For manual writing, select a section and inspect the linked plan and evidence.
2. Check the method and destination. A cited issue may suggest commenting on that issue. Expand destination content/recent status to confirm the original you intend.
3. Write the content or explicitly request an AI wording draft. Saving a draft prepares privately and does not send.
4. Review the actual account, destination and body in Preview before execution, then approve. Report/event handoffs still require final saving in their forms.
5. Record results and the next check afterwards. Later review drafts can use this action and confirmed observations as evidence.

### 2. Match evidence and verify the destination

1. Linked from evidence · No AI used means citations, links or exact mentions produced candidates. It is not an AI-confirmed decision or completed delivery.
2. Choose a method, then type a name, title, issue number or repository in Destination. Up to six suggestions appear without a Search button; More shows the next page. Before typing, only up to three exact evidence recommendations appear. Use arrow keys and Enter or click to select. The selection collapses to one line; Change preserves the title and body. Search uses no AI and distinguishes errors from no matches.
3. A cited source is not necessarily the place to publish. A comparison report does not automatically become the destination for an instruction. Decide whom you need to ask and why.
4. Inspect destination content/recent status for the body, assignees and available later comments. Freshness depends on connection type, access and read limits; check timestamps and errors.
5. If candidates are absent or ambiguous, select manually or open the original. Do not mistake a private personal record for an external destination.

### 3. Choose among AI action recommendations

1. AI recommendations prepares up to three alternatives from current evidence and permitted destinations. Adjust the direction toward an existing issue, direct inquiry, report/discussion or personal work, or add instructions.
2. Compare purpose, method, destination and reference numbers. Alternatives might ask in an existing issue, message a relevant researcher or discuss evidence below a report.
3. Recommendation batches are saved. After using one, expand the recommendations to create another draft. Continue draft opens existing wording; View result opens the completed action. Reuse consumes no AI allowance, and each action still requires review before posting.
4. New recommendations preserve earlier batches in a collapsed area. Check creation time and current destination context. Inaccessible destinations are unavailable; repeating an action requires explicitly choosing Create another action.
5. Use Change to correct an unsuitable destination manually. Multiple cited sources do not mean the same text should be sent everywhere.
6. AI recommendations and AI wording are separate requests: the first chooses a method/destination; the second drafts content for the selected destination.

### 4. Know where each method leads

1. GitHub comments post to an existing issue; new GitHub issues post to the selected repository. Chat messages go to the selected channel/direct conversation, and discussions go below the selected report/material. Verify access and destination.
2. PI check-in sharing prepares a report for its recipient in LUNA. Personal work/check-in is a private record with no external publication.
3. Manual/AI report creation and event creation prepare metadata and open the existing form. Recheck scope, period, dates, times, participants and audience there before saving.
4. Continue code work links the repository and task context. Opening a repository does not complete local changes, commits or pull requests.
5. For Notion or another supported publication, review the connection, actual destination and preview too. Publication does not imply overwriting the original or approving the research result.

### 5. Compose text, mentions, citations and attachments

1. Select a method and destination, then write the title and body. Optional AI wording is prepared for that destination and applied after review. Per-request instructions can specify questions and length.
2. Expand Edit & attach for Markdown headings, lists, tables, @mentions, [[document citations]], file selection, pasted images and drag/drop. Check upload limits and errors.

3. Select researchers from mention autocomplete. GitHub delivery uses explicitly linked GitHub IDs; a LUNA username is not assumed to be the same GitHub identity.

4. Check whether recipients can read cited documents. Links and attachments affect what is shared, so inspect the actual outgoing body and links.

5. Preview headings, lists, links and attachments before saving. If the source or destination changes, review a fresh preview instead of reusing earlier approval.

## 6. Drafts, approval and uncertain delivery

1. Save draft prepares an action. An AI draft marked ready means wording is ready, not that publication succeeded.

2. Review the sending account, destination, body, attachments and changes, then choose Confirm and execute. Complete any separate creation form in that form.

3. For success, inspect the result link and history. Sent does not mean read, answered, approved or research completed.

4. Execution needs checking or an unknown response can mean the external write succeeded without a receipt. Inspect the original and history before sending again.

5. If access, source or content changed, recheck the destination and use a new preview. Errors do not require deleting personal notes or original issues.

## 7. Inspect linked replies and later activity

1. Open action details → Results and next check → Check linked replies. Supported destinations are read for replies and activity after the action within bounded limits.

2. Directly linked replies and later activity on the same issue are labeled differently. A new issue comment is not necessarily an answer to your question; inspect content, author and time.

3. Background checking is supported but is not unlimited real-time monitoring. Recent actions are checked in bounded batches; access and connection limits can cause delays or omissions.

4. Check observation timestamps, errors and original links. No found reply does not prove inactivity. Inspect the original when needed.

5. Separate sending an instruction, receiving a reply and verifying an outcome. AI reviews preserve those distinctions and do not invent completion or approval.

## 8. Record personal results and share them

1. Personal work/check-in records something you do yourself. It is private and is not automatically published to the lab or a researcher.

2. Save directly verified results and next steps in Results and next check. A follow-up date schedules your own review; it is not an imposed recipient deadline or an automatic message schedule.
3. Share result creates a new sharing draft from the saved outcome. Choose an available chat, discussion, GitHub or report route, then preview and approve again. The original private record is not automatically published.
4. I finished my follow-up clears your follow-up candidate. It is different from external task completion, experimental success or PI approval.
5. Your results are self-reports; linked replies are observed evidence for AI review. They feed newly generated reviews/plans without automatically overwriting earlier saved text.

## 9. History, counts, cancellation and deletion

1. Use the number beside a plan or Follow-ups → View all to inspect actions. Event-history counts can include state changes and differ from confirmed execution counts beside plans.
2. Cancel a pre-execution draft with Cancel this action. This does not undo an existing GitHub comment, issue or chat. Check results first for processing or uncertain actions.
3. Expand Delete record in the list/details, review its irreversible scope and confirm. Draft, cancelled and posted records can be removed, but processing work needs outcome checks first.
4. Deleted personal actions are excluded from lists, detail views and AI evidence. Minimal receipts needed to prevent duplicates or authorize shared attachments may remain.
5. Deleting a personal execution record does not delete external issues, comments, chats, reports or original attachments. Use the original service's controls and permissions for those.

## 10. Recurrence, scheduling and GitHub access

1. Expand recurrence/scheduling only when needed. It starts disabled; recurring draft preparation and scheduled fixed-text delivery are different.
2. Scheduled delivery requires the complete actual destination/body preview and separate confirmation. Free-text instructions do not authorize arbitrary recurring external actions.
3. Changed sources, content or permission can stop a schedule or require renewed review. Check status/history and recurrence when cancelling or deleting.
4. In personal GitHub connection settings, connect only your account's required repositories. This does not grant access to other researchers' private scopes.
5. Disconnecting blocks later execution but leaves existing publications. Linking to repository work is distinct from actually committing code or opening a PR.

## 11. English display and translated drafts

1. In English mode, action lists, pickers and editors translate titles, destination labels, linked sources, recommendations and expanded context with the existing free translator on supported devices. Destination IDs and links stay unchanged and no personal AI allowance is used.
2. Korean Title, Content and Follow-up fields first show an English preview. Choose Edit original to revise the source or Use English draft to apply the translation to the field. Viewing a preview or saving a draft alone does not replace the original with English.
3. View original restores display text. It does not undo edits already applied with Use English draft. Late translations never overwrite text you are editing.
4. Mentions, document-citation syntax, link URLs, code, filenames, dates and Markdown markers are preserved. Review names, terminology and context before applying. The final execution review shows the actual saved text that will be submitted and is not automatically translated.
5. Translations of a saved action's title, body and follow-up criteria are reused as private copies. Changed source versions or access invalidate reuse; deleting the action clears its copies. Unsaved edits are translated only in the current page.
6. Choose Enable translation if the device needs preparation. Unsupported devices show a saved copy or the original, with no automatic switch to paid AI. Full pages and the action drawer in Today's research use the same behavior.

## 9. Report requests, responses and templates

Send, answer and review requests; manage templates, changes and deletion.

[Open this guide online](#)

### 1. Incoming requests, sent requests and received reports

1. Today's research, request lists and follow-up lists share the order Research suggestions → Received reports → Incoming requests → Sent requests → Follow-ups. Only the current tab is highlighted; swipe the row on narrow screens. Sent requests and Request a report appear when you have request permission. Incoming means you should answer; Sent means you requested it; Received reports contains submissions/shares addressed to you.
2. A red number counts requests whose body you have not yet opened. Select it to inspect new requests. Reading and submitting a response are separate states.
3. Opening a request marks it read, but it remains pending until submitted. Revision needed is separate. Without JavaScript, use I have read this request.
4. Use search, filters and pagination. Sent requests show the newest requests first; the response deadline and creation time are different values.
5. Request a report is available to platform administrators and currently authorized PIs within their permitted scope. Ordinary researchers can answer requests and use personal notes/actions.

### 2. Select recipients and exclusions

1. Choose individuals, a grant, a project or your full management scope. First verify the lab and scope you may manage.
2. Search individuals by name/username. Grant/project choices target participating researchers; expand Exclude recipients to remove unnecessary recipients.
3. Duplicate participation records are consolidated in the final recipient list. Review the actual names and count in the next step.
4. Loading a built-in or personal template preserves selected recipients, exclusions, period and deadline. It applies the request title, description and questions.
5. Participation and roles are rechecked at sending. Typing an unavailable person's name in the body does not add them as a recipient.

### 3. Write the request, period and deadline

1. Set the title, reporting start/end dates and response deadline. The period says which activities to describe; the deadline says when to answer. Check the displayed time zone.

2. Expand built-in/personal templates or write your own content. Read the Markdown preview and switch to Edit to change it.
3. Ask answerable questions, such as experiments performed, result links, blockers, decisions needed and next actions. Do not presuppose completion or submission.
4. AI drafting uses permitted context. Group drafts restrict context to avoid mixing one person's private notes into a group request. Personal allowances and access apply.
5. Review wording, recipients and period. Check the readable preview rather than sending raw JSON or internal identifiers. Saving a draft does not send it.

#### 4. Preview, send and understand notifications

1. After saving a draft, review the final recipient list/count, title, period, deadline and questions. Return to edit if anyone is missing or incorrect.
2. Explicit sending makes it appear in recipients' incoming requests. AI generation or preview alone does not deliver it.
3. Platform notifications and new-request indicators reflect the request state. Email/push depend on recipient preferences and delivery configuration; the sending screen does not prove email arrival.
4. Recheck a fresh preview after content or recipient changes. Do not reuse earlier approval when access or the request has changed.
5. Sending, reading, submitting a response and reviewing that response are separate stages. Inspect progress in Sent request details. The request instructions and questions are always visible.

#### 5. Answer, save and submit an incoming request

1. Open an incoming request and read the requester, period, deadline, questions and latest changes. Check the state first if the request was withdrawn.
2. Answer each question or import selected personal-note content. Share relevant outcomes and evidence rather than your whole private record.
3. When AI response drafting is available, still review actual work, links and uncertainty. Saved responses and AI drafts remain private before submission.
4. Review the submission preview and submit explicitly. Saving a personal goal or checking a daily box does not submit this answer.
5. For a revision request, read feedback, improve the response and resubmit. Inspect earlier content/change history and add the needed evidence.

## 6. Review reports, give feedback and remind

1. Use Received reports for submitted responses and PI check-in shares. Use an unread filter when available to prioritize unseen reports.
2. In sent request details, inspect each recipient's submitted, outstanding and review state. Read the body and evidence before giving feedback.
3. Below a requested report, Review and feedback shows all saved comments newest first. The feedback count excludes initial submissions and reviews without a comment. The top jump link and each comment's direct link preserve the report's current access checks.
4. Use Write new feedback below the saved comments. After saving, the server reads the saved comments again; Send feedback alone does not change workflow state. Submission and status history can be collapsed separately. Existing comments do not record edit timestamps, so edit status is not inferred. Reading or saving feedback itself has no AI cost.
5. For revisions, specify which answer needs what additional information. Review/read acknowledgment does not certify experimental validity or approve achievements.
6. Use the provided reminder control for outstanding recipients, subject to the 24-hour limit. Email/push delivery is different from an actual response.
7. Submissions are visible only to authorized recipients. Receiving private material from another scope does not authorize resharing it.

## 7. Edit, withdraw or permanently delete

1. Edit allowed title, description, period and deadline fields in Sent requests. Create a new request to change recipients or question structure after sending.
2. Existing responses and change history are preserved. Recipients may have answered an earlier version, so explain changes clearly.
3. Withdrawal stops the request while retaining responses/history. It differs from irreversible permanent deletion.
4. The author with current management permission can permanently delete draft, sent or withdrawn requests. Read the confirmation scope covering the request, linked responses, history and notifications.
5. Deletion does not undo personal check-ins, saved templates or AI usage. Review the scope carefully when responses need to be retained.

## 8. Create, edit, duplicate and delete templates

1. Open Template management beside built-in/personal template selection. Distinguish editing a template from the recipients/period of an in-progress request.

2. Edit a personal template's name, request title, description and questions, then save. Use a descriptive name such as Weekly experiment check.
3. Copy a built-in template into a personal one. Duplicate keeps a reusable original while creating a variation.
4. Confirm deletion to remove the template. It does not delete sent requests or responses. Template changes do not automatically update existing requests.
5. Period/deadline placeholders are filled when applied to the current request. Previously typed fixed dates are not automatically changed; check the preview. Saving and applying to the current request are separate choices.

## 9. Translation, privacy and troubleshooting

1. Korean requests/responses use existing automatic translation in English reading mode. Compare View original and inspect original text when editing.
2. The original can remain visible if browser translation is unsupported or not ready. Free on-device translation and generated AI translation are different features.
3. If Request a report is missing, check platform-admin or PI permission in the current scope. Ordinary project participation or changing language does not grant management rights.
4. If recipients look unexpected after loading a template, inspect selected conditions and the final list. Templates preserve recipients, exclusions, period and deadline while applying title, description and questions.
5. For deletion, withdrawal or edit conflicts, reopen the latest request. Send support the URL and error text, not private responses or tokens.

## 10. Research goals, literature and strategy

Paper goals, literature alerts, application advice, personal budgets and recurring strategy reviews.

Open this guide online

### 1. Set a personal paper goal

Open Today's research → Research suggestions in the Reports and actions row → + Paper goal. Select the lab and related projects, research question, venue and year, and your submission target date. Record available data, GPU time, collaboration constraints and PI/researcher perspective under Research time, resources and writing criteria. Goals are private and grant no additional source access. Use the Paper targets tab to set annual targets by category; legacy ambitious and other targets remain unclassified. Targets do not mean acceptance or publication; update the actual stage yourself.

### 2. Discover recent papers and receive application advice

Enter general English topic terms that may be shared publicly and enable Recent literature checks. Only those explicit terms go to OpenAlex, not unpublished research text automatically extracted as a search query. Checks enter a queue on a nominal 30-minute per-goal cadence and retrieve up to 25 results published in the past 14 days. Indexing delays, provider limits and queue load mean neither instant discovery nor exhaustive coverage is guaranteed. Discovered literature shows the last successful check and pending analysis count. Also enable automatic application advice/strategy to compare papers with your evidence and suggest an application, decisive experiment and limitations. A new-paper alert is created only when the analysis identifies concrete application value. Discovery alone uses no AI allowance.

### 3. Assess what can be done before your deadline

Enter a question such as "Suggest a paper and essential experiments feasible within ten days" and choose Strategy analysis. It uses bounded excerpts from currently accessible stored material, reports, personal notes and action results. With literature checks enabled it also searches earlier relevant work using the same public terms. Supported official venue/publisher HTML pages can supply a submission-guidance excerpt. Unread links and papers without full text remain explicit coverage limitations; registering a URL does not verify the rules. The report examines claims, prior-work differences, counter-evidence, decisive experiments, resources, critical path, fallback scope and potential overlap between papers. Feasible, conditional, rescope and evidence-needed are AI assessments, not acceptance probabilities. Indexes such as KCI and Scopus are not a single research-quality ranking.

## 4. Control analysis cost and notifications

Enable literature checks and automatic AI separately for each goal. Background analysis shares your personal allowance and must satisfy both a daily background cap and a reserve for manual use. Defaults are one background analysis and one reserved manual attempt per day. When allowance is unavailable, discovered papers remain stored for later analysis. A new analysis is marked as one attempt; reopening matching saved results needs no additional AI. Failed started attempts follow the existing quota policy. Quiet hours use your Today's research notification timezone. Daily digests combine ready suggestions after the selected hour; later arrivals wait until the next day. Email also requires existing email preferences and a configured sender; push requires a registered device and enabled notification category. Archiving a goal stops future discovery/automatic analysis while retaining its history.

## 5. Move from a report to plans and actions

Open Evidence and application from Research suggestions to inspect numbered citations and the source list. Act under Next actions opens the existing action editor; review and approve the destination and content before an external comment or message is posted. Add to today's plan appends the summary and actions without replacing existing notes. Adopt marks personal intent to use a suggestion, not paper acceptance or experimental completion. Use Already in progress, Show next week or Not interested to manage suggestions. Saved proposals inform AI consultation and plan/review drafts while remaining distinct from confirmed results. Subsequent execution results and new research records are reconsidered during later strategy analysis.

## 6. Recurring reviews, history and ChatGPT

Choose daily, every three days, weekly or fortnightly reviews. When automatic analysis is enabled, targets within seven days of the deadline are considered daily; changed internal evidence is reconsidered at least a day after the preceding analysis. New-paper analyses share the same personal cap. Earlier reports remain in the goal's version history. The ChatGPT connector can read saved goals/reports and prepare goal or personal-setting changes and a strategy analysis. Commit only after approving the exact review and cost. Poll the same goal to read analysis status instead of creating another request. External publication still needs the existing separate action preview and approval. If tools are unavailable, check reconnection and granted scopes.

## 7. Missing or delayed results

If a successful check finds nothing, broaden the public topic terms. Provider errors retain the last successful timestamp and retry on the next check. Papers waiting for analysis require automatic analysis enabled, sufficient cap/reserve and a working AI connection. Failed analysis retains the underlying evidence; correct the conditions and request analysis again. Changing project scope or source permissions may hide earlier reports; analyze the current scope again. Background discovery and analysis require the server worker to be running. The web page need not remain open.

## 8. Paper categories and counting

In Paper targets, select a year and View, then Add target. Choose KIISE conferences, SCI(E)/SCIE, Scopus, KCI or Custom and enter a submission count. Details hold the venue, official tier and edition, Q1 field, KCI candidate inclusion, main conference/workshop and reference link. These are user-defined criteria, not automatically verified accreditation. Keep one goal per manuscript and choose one primary annual counting category. Edit that same goal for resubmission, retaining its first submission date. Enter only confirmed submission, acceptance and publication dates; undated legacy stages are excluded from counts. Changing years preserves earlier settings and dates. Removing a target row removes only that annual target, not research records.

## 9. AI allowance examples

Under AI allowance, set the daily automatic cap and reserve for direct requests. With a total cap of 50, zero used, automatic cap five and reserve five, up to five automatic attempts are available. After 43 direct attempts, only two of the remaining seven can be automatic, leaving five. A total cap of five with a reserve of five stops automation. The reserve is not extra allowance and does not cap manual use. Automatic cap zero disables automation; reserve zero protects no allowance. Live calculations use the current usage snapshot and can change if you use AI elsewhere. The worker rechecks current quota and permissions.

## 10. Help and notification settings

Click or tap ? beside a setting for its meaning and examples. Keyboard Enter opens it and Escape closes it. In Notifications, choose delivery after analysis or a daily digest. Digest time uses your personal time zone and actual delivery follows the worker schedule. Quiet hours take precedence; equal start and end times disable the quiet interval. This does not pause analysis. Email also requires personal email preferences and a platform sender; it uses no extra AI. Save paper targets separately from AI/notification settings. Concurrent changes require a reload instead of silently overwriting saved settings.

## 11. Remove targets and record actual progress

New settings start with one target row. Add rows as needed; check Exclude this target and save to remove only the category. Uncheck before saving to undo. Existing targets are preserved. In Record progress, select the manuscript and enter owner-confirmed first submission, acceptance and publication dates. Each event counts once in its calendar year; a stage without a date is not counted. No external submission system automatically verifies these events. Reuse the same goal for resubmissions. Trashing a paper goal retains linked reports/code; restoration keeps automation off.

## 12. Shared costs and detailed dossiers

Only platform administrators control the shared public-analysis daily limit. Shared analysis of public papers/code has service cost but does not deduct personal AI attempts. Personalized strategy/code analysis uses personal quota. Expand Detailed report and downloads to set a maximum attempt budget for four saved report parts. Resume from saved parts after a budget stop or failure. Calls already started may still be charged after cancellation. Length and citation count follow available evidence; neither is padded with invented sources. ZIP includes Markdown, HTML, PDF, source JSON and BibTeX. Output examples are explicitly synthetic, not measured research, and cost no AI.

## 13. Progress and ChatGPT

Pages show saved results first and poll a small status response about every eight seconds while active or thirty seconds while idle, pausing in hidden tabs. Open View new results to refresh saved output. Polling uses no AI. ChatGPT uses guidance from consultant\_strategy\_list, status/dossier reads, and the same prepare/commit approval contract for progress, trash/restore and dossier start/pause/resume. If the host does not expose these tools, refresh connection metadata; server support alone does not prove host availability. Download links require a current platform login.

## 14. Connect existing context to a short question

Select grants/projects to include their purpose and constraints. The question editor supports mentions, document citations, Markdown tables and image/file attachments. Example: “Compare the model on page 6 with our project.” After saving, expand Research context to review the scope. Text extraction runs in the background after upload; analysis waits for pending inputs. Failed inputs can be retried. PDF text has extraction limits; the requested page or one question-relevant page is visually inspected. Check page numbers and omissions. Purpose, ideas, self-reports and validated results remain distinct; mentions do not send messages. Identify priority files/pages for large collections.

## 15. Progress, original requests, resuming and daily review

Staged strategy analysis saves the original question before running four parts: literature comparison; question/contributions/scope; model/experiments/implementation; schedule/integrated strategy. The default maximum is four personal AI attempts, one per stage. A lower cap pauses after saving completed parts. Waiting for another job, attachments or personal allowance does not lose your question. The compact status shows the target, current stage and saved parts. Expand request history for manual/automatic origin, original question, timestamps and errors. Attachment preparation and analysis completion are counted separately. Intermediate stages save concise, focused parts before final integration. Each stage has an AI response limit of up to eight minutes and continues after closing the page; this is not an estimated completion time. Timeouts pause; retry the saved request or resume from saved parts. Changed conditions require a new request. Download partial results as ZIP or Continue discussion without overwriting the original. Reading costs no platform AI; new analysis and consultation use personal allowance. Automatic failures back off and stop after three consecutive failures for the same goal revision. Daily AI review considers accessible, date-bounded goals and annual-target changes, analysis attempts/results, code proposals, working revisions and decisions. Missing historical snapshots or bounded coverage are not represented as complete collection.

### Collaborate from a shared report

Use Share and collaborate on a strategy report to review content, recipients and read or follow-up permissions. Report ID and immutable edition ID/version preserve lineage. Original attachments, private chats and GitHub access are not automatically shared. Open Shared research reports, then start your own discussion or code workspace. Share a saved result with the team by selecting an exact revision/checkpoint/code result and files to return to the author. Seen, review and adoption are separate human actions. Sharing, reading and exporting make no AI call; new generation uses the initiating user's allowance. Stable request IDs prevent duplicates. Revocation, expiry, withdrawal and project access changes affect subsequent access and continuation. Exported copies cannot be recalled. Notion queues a fixed edition under an authorized parent; unknown outcomes never auto-retry. Reconcile the existing page by ID. External researchers must first complete existing account and lab/project participation approval; member rollout is not automatically enabled.

### Active work, archives and deletion

The live strip shows queued/running requests only. Analysis archives separate attention, completed, cancelled/failed and trash; saved outputs remain downloadable. Cancel before moving active work to trash; restore never resumes automatically. Permanent deletion requires an exact scope review and current permission, and rejects references from later consultations/code. Original attachments and usage totals remain. Old encrypted backups are not immediately erased; reapply the latest deletion ledger after restoration. ChatGPT uses `consultant_strategy_prepare analysis_trash/analysis_restore/analysis_purge` and `commit`.

## Complete literature and report history

Goal settings precede research context. Paginate all literature/search observations and reports. Deduplicated literature and per-search snapshots are separate; legacy observations have unknown original search batches. Only the final analysis stage creates a final report; partial results remain on the request. Compare two saved reports without AI. ChatGPT `strategy_read` collection=`reports/literature/observations` with goal,page reads the same lists.

## Analysis continuity and ChatGPT

New strategy analyses default to prior reports, feedback, saved follow-up decisions and recent discussions for the same goal. Originals remain available; bounded excerpts and omissions are recorded for model input. Independent analysis excludes prior analysis context. Revoked evidence is never reused. Code consultation from a report, section or action retains source provenance. ChatGPT resumes the same request ID and checkpoints; completed work uses follow-up analysis. analysis with `status=true` returns small metadata. Review `strategy_prepare` acknowledge with kind/id/version before commit. If the host exposes an old limit or missing operations, refresh the connected tools and read the central policy.

## Synchronizing external discussion and code results

Each code workspace has one representative notice across preparation, AI and validation steps. Dismiss preserves failures and does not mark them successful. Link a strategy report under Sources and links; the report also links back to derived code. External results preserve the question, answer, files, provenance, versions and downloads. Ask ChatGPT or Codex to continue this LUNA workspace and save the answer and files to the same workspace. Only a ready storage receipt confirms synchronization. Conversations or temporary files never delivered to LUNA are not automatically imported. You can also paste an answer and select files in the platform. Conversation text is limited to 120 KB, files to 100 files and 8 MB total; retry with identical files and request identity to resume. Files are not executed; model and test claims are producer-reported. Revisions preserve previous results. Storage uses no new AI call but consumes storage and backup capacity. Encrypted rejected responses may be revalidated without AI for up to 24 hours; invalid responses remain failures. Execution/resume settings start collapsed.

## Evidence access and analysis recovery

Analysis detail shows unavailable evidence counts and Google reauthorization status. Open the linked connection, personally authorize Google, wait for collection, then check access again. Authorization alone never restores historical access. Reuse saved stages only after a new authorized collection confirms the same file, version and content. Changed content or revoked access keeps old results preserved but restricted. Review a new analysis with accessible evidence to create a separate four-stage request after reviewing the original question, budget and timeout. Blocked source text and old outputs are not copied; new calls consume the existing personal allowance. Preview/status reads never run AI. ChatGPT uses `strategy_read` analysis diagnosis and the exact `strategy_prepare` analysis\_recover / `strategy_commit` review contract.

## Enable or disable automation for a goal

Use the automation control at the top of a paper goal without deleting it. Conditions are collapsed by default. Periodic review offers None, daily, every three days, weekly and fortnightly. None disables periodic review only; new literature, changed evidence and deadlines within seven days are independent triggers. Disable automation to stop every automatic trigger. Literature discovery is separate and does not consume personal AI attempts. Disabling cancels unstated automatic requests; a started call may incur its existing cost and save a validated result before stopping. Manual analyses, reports, literature and history remain. Enabling or changing conditions schedules future work without replaying missed periods. Allowance waits, failure holds and running states are distinguished; the next review time is an estimate. Ask ChatGPT or AI consultation to disable automation for the goal. `consultant_strategy_read` returns automation and its revision; `prepare(operation=goal, values={automation_revision, automation:{automatic:false}})`, review the exact change, then commit updates the same policy. Actual host tool exposure must be verified separately. Reading and saving policy do not execute an AI analysis.

## Preserving accumulated literature and search history

Editing goal settings preserves discovered papers, analysis links and observations. Titles and deadlines do not change the search revision; changing search terms does. Accumulated literature is the default, with filters for current, earlier or unknown settings, recent discovery, analysis availability and relevance review, paginated by 20. Unknown historical search terms are never inferred. Prioritize, select by relevance and exclude apply to future analyses without deleting previous results. Reassess relevance after changing the question or research scope. Reads and selection saves spend no AI allowance; automatic analysis consent is separate. `consultant_strategy_read(goal)` returns the same library in ChatGPT and AI consultation. `prepare(operation=goal, values={literature:{operation:list,filter:all,page:1}})` is read-only. For changes, use literature operation=select, paper, inclusion(auto/include/exclude) and revision, review, then commit. Search observations retain the actual terms and settings.

## 3. Sharing, Notion and printing

Use the top share icon for team sharing or Notion publication. Find recipients by name and use ? for field help. For Notion, select an allowed parent and review the body before queuing publication. Configure a connection if none is available. Notion audiences are independent. The export icon provides full-body print, PDF and Markdown. An uncertain publication is reconciled, never blindly reposted.

## 4. ChatGPT and AI consultation

Connected tools support reading a goal workflow and recording notes or imported results. Use the lifecycle envelope inside `consultant_thread_prepare` save, followed by reviewed commit for changes. Only transmitted content is saved. Reads and record writes do not invoke extra AI or remote experiments; new consultation follows existing usage policy.

## 4. Continue across AI surfaces

1. Use  beside the send button in either a new or follow-up request to choose: Answer in this work, Continue in AI chat, or Continue in ChatGPT.

2. Selecting a destination runs no AI. Answer in this work sends the request. Continue in AI chat prepares your final question, selected evidence and files in the consultation panel; review and send there to start.
3. Continue in ChatGPT prepares handoff text. Copy it into ChatGPT with LUNA connected. Copying alone neither transfers nor saves anything externally.
4. Reuse the same work and conversation for the same purpose. New work is created on actual send; switching a prepared request's destination must not run it twice.
5. Only delivered and saved content and files are preserved. Check the receipt's work ID, files and download links. Choosing a destination does not broaden private conversation access or sharing.

## 7. Core materials and AI drafts

1. Search and pin important sources in Materials & results. Pinning and unpinning do not copy, delete or change source access.
2. Private pins enter only your own AI context; team sources still require current source permissions. Attach relevant evidence and choose Ask AI to save a request on the same work.
3. Choose Save note to preserve a private memo without AI; it does not replace the current result. Waiting, running and failure reflect actual request state; saved stages remain preserved.
4. AI requests follow administrator policy and personal limits. Reading, saving notes and pinning make no new generation call.
5. Pins and AI drafts are not factual verification or researcher approval.

## 10. Find earlier analyses and results

1. Existing goals, code, reports and discussions retain their original IDs and permissions while being linked into one work list. Open work or search Materials & results for originals and downloads.
2. The current revision and latest saved checkpoint represent a conversation; earlier versions remain in detailed history. Distinct reports are not merged by matching titles.
3. If evidence is inaccessible, the owner may see only a preserved-result notice. Restore current connection and permissions before opening content or downloads.
4. Never infer or claim to have read blocked content. A classification-only change with identical content, version, connection and current access does not require regenerating saved results.

## 11. One work list and research map

1. Choose List or Research map in Research work. Selecting work opens the right panel. Use ↗ or Ctrl/⌘-click to open the same work on its own page with a stable URL. The default map groups results inside their owning work; Expand outputs reveals exact report and code-result cards.

2. Renaming work changes its display name without rewriting source titles, content, versions or evidence, or duplicating work. Distinct works are never merged by title.
3. Solid workflow/derivation links and dashed reference/validation links are both visible by default. Multiple origins, references and sections retain their exact report IDs and versions even when output cards are folded.
4. Counts distinguish the works, outputs and materials currently shown, not the total saved collection. Use the next work range or a focused neighborhood to navigate.
5. Renaming or reordering the list does not reposition the same works and relationships. Goal settings, literature, strategy analysis and specialized code tools remain available from that work's settings or detail screen.
6. Source changes require deliberate evidence review and refresh. Disconnecting does not delete the source or earlier history.
7. Maps, search and opening a work URL call no AI and create no work or conversation. The panel and separate page use the same work, current permissions, conversation, results, files and history.
8. While drafting, use ↗ to open a new tab. Unsent text and selected files remain in the original tab. Open in research space returns to that work. Use the original tab to resume its search, map position and draft.

## 12. Trash and a new research cycle

1. Review the target and effects before moving work or a space to Trash. Saved conversations, outputs, attachments and external originals remain preserved.
2. A new cycle archives earlier work while retaining materials and members. Restore from Trash or previous cycles without automatically restoring work individually trashed earlier.
3. Related local pending or running requests stop, but a call already sent to an external provider cannot be guaranteed recalled. Restoration itself queues no AI.
4. Original automation policies remain preserved, so an enabled policy may execute at its next scheduled time after restoration.

## 13. Navigate quickly and inspect history

1. Switch between List and Map and sort by recent activity, creation, execution state, attention or title. Changing work loads its detail in place and retains the draft question for each work.
2. Open Records and filter All, Conversation, Results and versions, Executions or Links. An answer and its result are grouped.
3. A conversation started from an earlier report or selected passage retains its original thread. Load earlier records on demand.

4. No execution requested, generating and needs attention describe execution separately from saved outputs. Downloads require actually preserved files.

## 14. Save naturally from a conversation

1. Ask the AI consultation or ChatGPT: "Save this conversation and its results to the current research work." LUNA prioritizes an explicit work, the bound conversation and a previous save receipt. Choose by title when several targets match, or preserve privately in the personal inbox before selecting a destination.
2. New work requires an explicit request; retrying the same request does not duplicate it. Review the destination, content and files before saving.
3. The personal inbox supports downloads and linking to existing work. Saving text makes no additional platform model call.
4. Files can be uploaded in resumable chunks up to 8 MB per bundle; larger materials may be saved as accessible source references. A source link is not a preserved binary copy.
5. Check the completed save receipt: uploading means only received parts are preserved. External messages and files that were never delivered are not automatically collected.
6. Team sharing and external publication retain separate scope reviews.

## 15. Full-width map and images

1. The map uses the full work area. Derivation and dependency relations guide its left-to-right layout, including independent work and loops.
2. Select a node for detail and return to the same map position. More offers PNG, JPEG and SVG exports of the currently loaded range, not a claim to cover every record.
3. No AI call is made.

## 16. Attach files and use code results

1. Use Add materials for files, images and platform evidence; paste links into your question. Upload up to 50 files totaling 100 MB through resumable chunks.
2. Originals are preserved while AI excerpts are prepared in the background. Preparation is not a full read of every original; extraction failures and text or image limits are shown separately.
3. Actually completed ChatGPT file transfers can be attached with the same permission checks. Missing code output has no download; execution history shows the failure.
4. Import preserved text code or a validated ZIP as a code changeset and inspect each diff. ZIPs are never executed or automatically installed.

5. Review repository, branch, baseline commit and changes before approving GitHub publication. Reconcile uncertain responses to avoid duplicate posts.

## 17. Keep existing work in My saved items

1. Star existing work, conversations or results to keep a shortcut in My saved items. This differs from a preserved copy of delivered conversation content and files.
2. Shortcuts reopen the same work and version; source permission revocation blocks content access. Web, AI consultation and ChatGPT use the same bookmarks and save receipts.

## 18. One library and pinned materials

1. Use All, Favorites and Saved conversations and files in one personal library. Old favorite URLs and folders remain valid.
2. Unstarring does not delete preserved originals. Materials and results initially shows pinned items; search includes all accessible materials.
3. Toggle the same pin to unpin and clear search to return to pins. Pinning does not expand source permissions.

## 19. Private work and team discussion

1. Click Only me on a work item to manage sharing. Review the exact body and selected files before sharing.
2. Earlier private discussions and personal memory are not shared automatically. Continue in team view to accumulate attributed discussion under the same canonical research work.
3. Sharing withdrawal, membership and source access are rechecked; private originals remain preserved. ChatGPT uses the same work ID and current permissions and reflects only delivered content.

## 20. Add schedules and export maps

1. Use Add to choose an event, GitHub task or report request and then create or link existing. Creation opens the standard form with optional AI drafting.
2. Linking does not duplicate the original or create an external item. Select one image format before downloading a map.
3. Dashed reference and validation relations identify evidence or verification links, separately from derivation order.

## 24. Refine a request through conversation

1. Write a few sentences in New research or Continue this work, then select sources and files in Add materials. Put links in the text. Both composers use the same refinement, attachment and destination controls.
2. Refine request opens a conversation on the right. It checks the current request, work and selected evidence, suggests a research kind and identifies missing conditions. It reuses known conditions and asks one or two important questions at a time.
3. A choice fills the reply box. Add or edit details, then choose Send reply. You may leave a condition undecided or ask for a recommendation. Repository candidates are authorized connection metadata, not proof of write permission or server availability.
4. Preview the prepared request and conditions, distinguishing your input, existing evidence, AI suggestions and unknowns. The composer changes only when you choose Apply to request. Undo restores your earlier input.
5. Applying does not execute research. Edit the result if needed, choose this work, AI consultation or ChatGPT from the destination menu, then send. You may also skip refinement and send directly.
6. Closing and reopening preserves the saved conversation. On mobile, closing returns to the original composer. An older response never automatically overwrites text you changed.
7. Starting refinement and sending each additional reply request one personal AI attempt each. Page reads, saved conversation reads and applying use no AI. Closing the panel does not cancel an already accepted AI job.
8. Administrator policy, remaining allowance, provider and time limits apply. Failures retain input and conversation without automatic paid retries. Up to 50 files totaling 100 MB are supported; uploading is distinct from analyzing all original content.
9. Selecting a saved suggestion only fills the request and evidence; Save memo stores a record without AI. Team requests use only currently authorized shared evidence.

## 25. Continue request drafts and results in ChatGPT

1. You can hand off a rough request without finishing refinement on the website. Choose Continue in ChatGPT from the destination menu for either a new request or a follow-up.
2. The handoff includes the latest composer text, selected sources, prepared attachments and existing work, conversation and result versions. An applied refinement also carries the same request session and accepted revision. Preparing and copying the handoff calls no AI, creates no work and transfers nothing externally.
3. Paste the text into ChatGPT with LUNA connected and check that the host exposes the needed tools. Ask: "First clarify the request by asking for missing conditions; proceed after we settle it."

4. Preserve the same request session, work and evidence versions when saving a request refined in ChatGPT. Saving externally prepared requests and results does not rerun platform AI. Draft storage, applying a request and executing research remain separate actions.
5. Ask: “Save the discussion delivered so far and the generated files to this work, then show download links.” The platform preserves only conversation content and files actually delivered through its tools, not the entire external chat automatically.
6. A filename or temporary link is not proof of preservation. Check the receipt destination, file list, downloads and readiness. Use web upload if the host cannot deliver files.
7. Confirm a destination for new work and retain the existing work, conversation and exact result version for follow-ups. Continue through the stable work link; include the exact result or version link when reviewing a specific output. Opening a link does not save drafts, external conversations or files. Private conversations and evidence do not become team-shared automatically. Publication and paid GPU execution have separate scopes.

### 33. Automatic recommendations and usage

1. Open space settings in Research team to turn automatic research recommendations on or off. Existing and new spaces default to on, preparing suggestions from existing reports and consultations.
2. Opening pages, reading suggestions, searching and sorting make no new AI call. Relevant saved results refresh recommendations in the background; the first view shows saved suggestions.
3. Additional automatic AI enrichment defaults to zero attempts. Set a daily cap if needed; the space owner’s allowance and administrator policy also apply. Changing settings never resets consumed usage.
4. Turning off stops subsequent automatic recommendation generation while preserving existing suggestions and history. Pending work rechecks the setting and permissions immediately before execution.
5. Private consultation or source recommendations are not automatically shared with other members. Changed or revoked evidence requires revalidation. Preparing a suggestion does not execute research, prove success or constitute human review.

### 34. Compact cards and work search

1. Open titles in the Projects and Research spaces sections of a grant or project. The directory uses the same cards without repeating a parent already shown.
2. Work and result counts use currently accessible entries. A work’s conversation, report and execution history are not counted as separate work. A + marks a bounded count.
3. Choose List or Map and type to search Research work. Search waits for composition to finish; Enter searches immediately. Use the compact sort menu for recent activity, creation, status and other criteria.

4. Searching and changing views preserve the draft and selected work. Renaming work or adjusting map relations retains the same work ID and history.

### 35. Research kinds, requests and results

1. A research kind describes the activity: planning, literature review, model design, experiment design, code implementation, training, validation, reproducibility, analysis, discussion, writing, reporting or scheduling. Reports, code files and literature tables are result formats.

2. AI proposes a kind during refinement. Direct requests receive an initial classification; use Change research kind in the work menu to correct it. Manual choices are never overwritten automatically.

3. Every work uses Request, Results and Continue this work, in that order. Expand long requests and open or download results. Find initial requests, conversation, execution and earlier versions in History under the work menu.

4. Changing kind preserves sources, results and permissions; it does not authorize server execution or GitHub writes. Distinguish a training plan from actual training and a validation plan from completed validation.

### 36. Continue earlier goals and literature

1. Earlier paper-goal settings and discovered literature can be organized as a separate preparation work. Read its research objective and conditions as the request, and open the literature bundle as its result.

2. The literature result includes search and observation history, selection state and sources. Open or download it, distinguishing reorganized existing records from new AI analysis.

3. Earlier strategy and detailed-report IDs, conversations, reports, citations and downloads remain intact. Linking preparation as the logical starting point never invents historical runs or conversations.

4. Earlier settings and literature URLs lead to the preparation work. Its automatic execution settings retain the earlier policy; reorganization does not resume stopped AI or incur additional AI charges.

5. Renaming work never creates another map node for the same work. Work and result artifacts remain distinct; expand result versions when needed.

# 11. Shared literature, research code and experiments

Shared literature, repository-aware coding, ablations, CPU checks and draft PRs.

[Open this guide online](#)

## 1. Shared literature and private research

Public papers are linked to a shared database by DOI or OpenAlex identifier. Matching queries reuse cached discovery and general analysis of the same public source is reusable. Private goals, query histories, repository code and personal applicability advice are not shared with other users. Existing paper IDs and report citations are preserved. Only platform administrators can set the shared-analysis budget. Its default daily cap of zero disables extra shared AI analysis while search caching still works. When enabled, public-only analysis uses the separate platform budget, never the first discoverer's personal allowance.

## 2. Connect an idea to a code workspace

Open Research goals and suggestions → Code work, or Apply to code on a goal or strategy report. A GitHub code action also offers Continue with code analysis and implementation. Choose your authorized GitHub repository and optional paper goal, then describe the idea. A blank base branch uses the repository default. Optionally supply the paper's public GitHub repository URL. Choose Read repository / check again first. This read has no AI charge and requires actual repository read access.

## 3. Inspection coverage and change watching

Expand Repository and inspection coverage to see the exact commit, branch, inspected/omitted files, functions, classes and license. Each read is bounded to 24 text files and 85KB; uninspected code and runtime remain unknown. Enter missing paths one per line under Priority file paths and read again. Public upstream code is checked anonymously for public access. Optional six-hour watching checks for changes. Automatic new recommendations additionally require the explicit AI toggle and obey the personal background cap and manual reserve. It never automatically publishes code or PRs.

## 4. Compare alternatives and generate implementation or ablation code

AI code recommendations use one personal attempt to propose up to three approaches. Compare files, integration points, validation and limitations, then choose Implementation draft or Ablation code; generation is a separate attempt. Saved recommendations and changes let you reopen earlier alternatives without regenerating recommendations. Matching input and commit reuse saved output. Inspect file diffs and optionally edit complete file contents into a new version. Review upstream licensing, attribution, compatibility and evidence for any research claim.

## 5. CPU checks and actual experiments differ

CPU checks run Python syntax checks and unittest discovery in tests inside an environment separated from the web service. No network, GPU or production files are available; limits are 65 seconds, 384MB and ten checks per day. No dependencies are installed, so dependency-heavy tests may fail. Passed means only those checks passed, not improved accuracy, reproducibility or a paper result. Inspect the details and revise the patch or use the required environment. CPU checks are separate from AI allowance and do not publish code.

## 6. Run ablations locally or in your cloud environment

Download experiment package exports the patch, base commit, conditions, seeds, commands and runner. Review README and code, then use a clean separate clone at the exact base commit. Apply with `python luna_experiments.py --repo PATH --apply --approve`; existing changes stop the operation. Prepare libraries, data and GPU yourself. Run with `--run --approve --seconds 600` to set the total time limit. `LUNA_SEED` and `LUNA_RESULT_PATH` communicate the seed and output location. The runner does not install, commit, push or provision paid servers. Inspect the result JSON before importing it into the matching artifact. Imported results are self-reported, not independent platform reproduction.

## 7. Approve changes and create a draft PR

Under Apply as a GitHub draft PR, acknowledge review of changes, licenses and repository visibility, then open the preview. Review the repository, base commit, file changes and check results before confirming. Multiple files are committed together on a separate `luna/research-` branch and a draft PR is created. Main is not updated and nothing is merged. A changed base branch requires a fresh read and review. If a response is uncertain, use Check publication result instead of reposting. GitHub write permissions and your personal execution connection or an authorized management connection are required.

## 8. Continue through strategy, daily reviews and ChatGPT

Saved implementations, CPU checks, self-reported experiments and PR links inform later strategy and AI plan/review drafts. Code generation, test passing, publication and research success remain distinct. `consultant_code_list/read` retrieve the same workspaces, alternatives and patches without new platform AI. `consultant_code_prepare/commit` provide exact reviewed creation, scanning, recommendations, implementation, ablation, CPU checks, result import and draft PR operations. PR publication additionally requires research-action write scope. Private analysis and personalized implementation consume personal AI allowance; only generic public-source analysis is shared. If tools are missing, refresh the connection and check granted scopes.

## 9. From an idea to code, downloads and publication

Use Implement in code beside a report action to retain that specific idea and its references. Select an authorized repository/base branch, then read, recommend and implement. Other saved alternatives remain available. Generation requests architecture/API/usage documentation and meaningful toy tests, but output is unverified until executed. The ZIP includes changed files, patch, explanation and checks; individual files can also be downloaded. Default publication creates a new branch and draft PR. Direct commit to selected base branch requires review and approval of the exact branch, baseline commit and diff. No force update; remote drift or protection requires review. ChatGPT consultant\_code\_prepare uses the same mode and report/idea linkage.

## 10. Multiple workspaces and input help

Create a workspace with + Code work. Multiple independent workspaces can share a repository or paper goal and retain separate recommendation, implementation and validation histories. The ? beside each of six fields explains its meaning and examples. A linked goal contributes project/grant purpose, explicit document citations and prepared attachments. Wait for pending extraction before analysis. The base branch is the reading baseline; creating a workspace does not commit. Lists are paginated; open one workspace for details. Shared navigation keeps Research development and Communication/follow-up in the same order. The integrated ChatGPT guide also covers search, reports, events, messages, actions and help.

## 11. Editing ideas, attachments and when commits happen

Expand Edit and attach below Idea / implementation to use headings, bold, tables, @mentions, [[document citations]], images, files and preview. Select a GitHub connection first. Files stay private to the workspace; supported content is extracted in a background job. Limits are 1.5MB per image, 20MB per file and 20 attachments; configured private Google Drive storage is required. Check original access and preparation status before analysis. PDFs are not entirely converted into AI images; specify relevant page or figure numbers. Retry failed preparation without AI usage. Unsupported extraction formats such as ZIP remain downloadable files, not analyzed content. Edit existing workspaces under Repository and read coverage. Create workspace only saves a platform record: no commit, PR or mention notification. Apply to GitHub requires a separate review of changes, branch and commit/PR title followed by approval. The workspace title is currently used for both titles. ChatGPT consultant\_code\_read exposes the same attachment preparation states and platform AI uses the same prepared evidence. Uploading a file to ChatGPT does not automatically store it in LUNA; attach it in the platform code editor and continue using its workspace ID.

## Repository consultation, isolated checks and integration

In GitHub code consultation choose architecture/paper mapping, README, function/class documentation or commit/research changes. The inspected file coverage and up to 20 commit messages are explicit; execution is never inferred. ZIPs include inspected source plus changes under standalone/, INTEGRATION.md, patches and check receipts. Run toy examples and tests against the real generated API in an isolated environment before integrating on a separate branch at the pinned baseline. Zero tests or missing dependencies do not establish verification.

## Detailed filters and research notifications

Detailed filters search authorized users, grants and projects even without recorded usage. Select all translation or individual features including help. Unassigned means missing attribution; unmeasured tokens are not zero. The top area shows active work, issues and unseen results. Mark as seen acknowledges a notification, not scientific approval. Cancelled and reviewed work remains in the goal/code history and archive. Saved materials are preserved.

## Analysis continuity and ChatGPT

New strategy analyses default to prior reports, feedback, saved follow-up decisions and recent discussions for the same goal. Originals remain available; bounded excerpts and omissions are recorded for model input. Independent analysis excludes prior analysis context. Revoked evidence is never reused. Code consultation from a report, section or action retains source provenance. ChatGPT resumes the same request ID and checkpoints; completed work uses follow-up analysis. analysis with status=true returns small metadata. Review strategy\_prepare acknowledge with kind/id/version before commit. If the host exposes an old limit or missing operations, refresh the connected tools and read the central policy.

## Synchronizing external discussion and code results

Each code workspace has one representative notice across preparation, AI and validation steps. Dismiss preserves failures and does not mark them successful. Link a strategy report under Sources and links; the report also links back to derived code. External results preserve the question, answer, files, provenance, versions and downloads. Ask ChatGPT or Codex to continue this LUNA workspace and save the answer and files to the same workspace. Only a ready storage receipt confirms synchronization. Conversations or temporary files never delivered to LUNA are not automatically imported. You can also paste an answer and select files in the platform. Conversation text is limited to 120 KB, files to 100 files and 8 MB total; retry with identical files and request identity to resume. Files are not executed; model and test claims are producer-reported. Revisions preserve previous results. Storage uses no new AI call but consumes storage and backup capacity. Encrypted rejected responses may be revalidated without AI for up to 24 hours; invalid responses remain failures. Execution/resume settings start collapsed.

## Start repository consultation

Choose a project, repository and purpose using Start GitHub repository consultation on the code landing page. Continue an existing analysis where available. Consultation is expanded on the detail page and requires a completed repository read. Connection/permission problems link to connection settings. The current bound of 24 text files/85KB does not establish whole-repository verification.

### 4. Continue across AI surfaces

1. Use  beside the send button in either a new or follow-up request to choose: Answer in this work, Continue in AI chat, or Continue in ChatGPT.
2. Selecting a destination runs no AI. Answer in this work sends the request. Continue in AI chat prepares your final question, selected evidence and files in the consultation panel; review and send there to start.
3. Continue in ChatGPT prepares handoff text. Copy it into ChatGPT with LUNA connected. Copying alone neither transfers nor saves anything externally.
4. Reuse the same work and conversation for the same purpose. New work is created on actual send; switching a prepared request's destination must not run it twice.
5. Only delivered and saved content and files are preserved. Check the receipt's work ID, files and download links. Choosing a destination does not broaden private conversation access or sharing.

### 7. Core materials and AI drafts

1. Search and pin important sources in Materials & results. Pinning and unpinning do not copy, delete or change source access.
2. Private pins enter only your own AI context; team sources still require current source permissions. Attach relevant evidence and choose Ask AI to save a request on the same work.
3. Choose Save note to preserve a private memo without AI; it does not replace the current result. Waiting, running and failure reflect actual request state; saved stages remain preserved.
4. AI requests follow administrator policy and personal limits. Reading, saving notes and pinning make no new generation call.
5. Pins and AI drafts are not factual verification or researcher approval.

### 10. Find earlier analyses and results

1. Existing goals, code, reports and discussions retain their original IDs and permissions while being linked into one work list. Open work or search Materials & results for originals and downloads.
2. The current revision and latest saved checkpoint represent a conversation; earlier versions remain in detailed history. Distinct reports are not merged by matching titles.

3. If evidence is inaccessible, the owner may see only a preserved-result notice. Restore current connection and permissions before opening content or downloads.
4. Never infer or claim to have read blocked content. A classification-only change with identical content, version, connection and current access does not require regenerating saved results.

## 11. One work list and research map

1. Choose List or Research map in Research work. Selecting work opens the right panel. Use ↗ or Ctrl/⌘-click to open the same work on its own page with a stable URL. The default map groups results inside their owning work; Expand outputs reveals exact report and code-result cards.
2. Renaming work changes its display name without rewriting source titles, content, versions or evidence, or duplicating work. Distinct works are never merged by title.
3. Solid workflow/derivation links and dashed reference/validation links are both visible by default. Multiple origins, references and sections retain their exact report IDs and versions even when output cards are folded.
4. Counts distinguish the works, outputs and materials currently shown, not the total saved collection. Use the next work range or a focused neighborhood to navigate.
5. Renaming or reordering the list does not reposition the same works and relationships. Goal settings, literature, strategy analysis and specialized code tools remain available from that work's settings or detail screen.
6. Source changes require deliberate evidence review and refresh. Disconnecting does not delete the source or earlier history.
7. Maps, search and opening a work URL call no AI and create no work or conversation. The panel and separate page use the same work, current permissions, conversation, results, files and history.
8. While drafting, use ↗ to open a new tab. Unsent text and selected files remain in the original tab. Open in research space returns to that work. Use the original tab to resume its search, map position and draft.

## 12. Trash and a new research cycle

1. Review the target and effects before moving work or a space to Trash. Saved conversations, outputs, attachments and external originals remain preserved.
2. A new cycle archives earlier work while retaining materials and members. Restore from Trash or previous cycles without automatically restoring work individually trashed earlier.
3. Related local pending or running requests stop, but a call already sent to an external provider cannot be guaranteed recalled. Restoration itself queues no AI.

4. Original automation policies remain preserved, so an enabled policy may execute at its next scheduled time after restoration.

### 13. Navigate quickly and inspect history

1. Switch between List and Map and sort by recent activity, creation, execution state, attention or title. Changing work loads its detail in place and retains the draft question for each work.

2. Open Records and filter All, Conversation, Results and versions, Executions or Links. An answer and its result are grouped.

3. A conversation started from an earlier report or selected passage retains its original thread. Load earlier records on demand.

4. No execution requested, generating and needs attention describe execution separately from saved outputs. Downloads require actually preserved files.

### 14. Save naturally from a conversation

1. Ask the AI consultation or ChatGPT: “Save this conversation and its results to the current research work.” LUNA prioritizes an explicit work, the bound conversation and a previous save receipt. Choose by title when several targets match, or preserve privately in the personal inbox before selecting a destination.

2. New work requires an explicit request; retrying the same request does not duplicate it. Review the destination, content and files before saving.

3. The personal inbox supports downloads and linking to existing work. Saving text makes no additional platform model call.

4. Files can be uploaded in resumable chunks up to 8 MB per bundle; larger materials may be saved as accessible source references. A source link is not a preserved binary copy.

5. Check the completed save receipt: uploading means only received parts are preserved. External messages and files that were never delivered are not automatically collected.

6. Team sharing and external publication retain separate scope reviews.

### 15. Full-width map and images

1. The map uses the full work area. Derivation and dependency relations guide its left-to-right layout, including independent work and loops.

2. Select a node for detail and return to the same map position. More offers PNG, JPEG and SVG exports of the currently loaded range, not a claim to cover every record.

3. No AI call is made.

## 16. Attach files and use code results

1. Use Add materials for files, images and platform evidence; paste links into your question. Upload up to 50 files totaling 100 MB through resumable chunks.
2. Originals are preserved while AI excerpts are prepared in the background. Preparation is not a full read of every original; extraction failures and text or image limits are shown separately.
3. Actually completed ChatGPT file transfers can be attached with the same permission checks. Missing code output has no download; execution history shows the failure.
4. Import preserved text code or a validated ZIP as a code changeset and inspect each diff. ZIPs are never executed or automatically installed.
5. Review repository, branch, baseline commit and changes before approving GitHub publication. Reconcile uncertain responses to avoid duplicate posts.

## 17. Keep existing work in My saved items

1. Star existing work, conversations or results to keep a shortcut in My saved items. This differs from a preserved copy of delivered conversation content and files.
2. Shortcuts reopen the same work and version; source permission revocation blocks content access. Web, AI consultation and ChatGPT use the same bookmarks and save receipts.

## 18. One library and pinned materials

1. Use All, Favorites and Saved conversations and files in one personal library. Old favorite URLs and folders remain valid.
2. Unstarring does not delete preserved originals. Materials and results initially shows pinned items; search includes all accessible materials.
3. Toggle the same pin to unpin and clear search to return to pins. Pinning does not expand source permissions.

## 19. Private work and team discussion

1. Click Only me on a work item to manage sharing. Review the exact body and selected files before sharing.
2. Earlier private discussions and personal memory are not shared automatically. Continue in team view to accumulate attributed discussion under the same canonical research work.
3. Sharing withdrawal, membership and source access are rechecked; private originals remain preserved. ChatGPT uses the same work ID and current permissions and reflects only delivered content.

## 20. Add schedules and export maps

1. Use Add to choose an event, GitHub task or report request and then create or link existing. Creation opens the standard form with optional AI drafting.
2. Linking does not duplicate the original or create an external item. Select one image format before downloading a map.
3. Dashed reference and validation relations identify evidence or verification links, separately from derivation order.

## 24. Refine a request through conversation

1. Write a few sentences in New research or Continue this work, then select sources and files in Add materials. Put links in the text. Both composers use the same refinement, attachment and destination controls.
2. Refine request opens a conversation on the right. It checks the current request, work and selected evidence, suggests a research kind and identifies missing conditions. It reuses known conditions and asks one or two important questions at a time.
3. A choice fills the reply box. Add or edit details, then choose Send reply. You may leave a condition undecided or ask for a recommendation. Repository candidates are authorized connection metadata, not proof of write permission or server availability.
4. Preview the prepared request and conditions, distinguishing your input, existing evidence, AI suggestions and unknowns. The composer changes only when you choose Apply to request. Undo restores your earlier input.
5. Applying does not execute research. Edit the result if needed, choose this work, AI consultation or ChatGPT from the destination menu, then send. You may also skip refinement and send directly.
6. Closing and reopening preserves the saved conversation. On mobile, closing returns to the original composer. An older response never automatically overwrites text you changed.
7. Starting refinement and sending each additional reply request one personal AI attempt each. Page reads, saved conversation reads and applying use no AI. Closing the panel does not cancel an already accepted AI job.
8. Administrator policy, remaining allowance, provider and time limits apply. Failures retain input and conversation without automatic paid retries. Up to 50 files totaling 100 MB are supported; uploading is distinct from analyzing all original content.
9. Selecting a saved suggestion only fills the request and evidence; Save memo stores a record without AI. Team requests use only currently authorized shared evidence.

## 25. Continue request drafts and results in ChatGPT

1. You can hand off a rough request without finishing refinement on the website. Choose Continue in ChatGPT from the destination menu for either a new request or a follow-up.
2. The handoff includes the latest composer text, selected sources, prepared attachments and existing work, conversation and result versions. An applied refinement also carries the same request session and accepted revision. Preparing and copying the handoff calls no AI, creates no work and transfers nothing externally.
3. Paste the text into ChatGPT with LUNA connected and check that the host exposes the needed tools. Ask: "First clarify the request by asking for missing conditions; proceed after we settle it."
4. Preserve the same request session, work and evidence versions when saving a request refined in ChatGPT. Saving externally prepared requests and results does not rerun platform AI. Draft storage, applying a request and executing research remain separate actions.
5. Ask: "Save the discussion delivered so far and the generated files to this work, then show download links." The platform preserves only conversation content and files actually delivered through its tools, not the entire external chat automatically.
6. A filename or temporary link is not proof of preservation. Check the receipt destination, file list, downloads and readiness. Use web upload if the host cannot deliver files.
7. Confirm a destination for new work and retain the existing work, conversation and exact result version for follow-ups. Continue through the stable work link; include the exact result or version link when reviewing a specific output. Opening a link does not save drafts, external conversations or files. Private conversations and evidence do not become team-shared automatically. Publication and paid GPU execution have separate scopes.

## 33. Automatic recommendations and usage

1. Open space settings in Research team to turn automatic research recommendations on or off. Existing and new spaces default to on, preparing suggestions from existing reports and consultations.
2. Opening pages, reading suggestions, searching and sorting make no new AI call. Relevant saved results refresh recommendations in the background; the first view shows saved suggestions.
3. Additional automatic AI enrichment defaults to zero attempts. Set a daily cap if needed; the space owner's allowance and administrator policy also apply. Changing settings never resets consumed usage.
4. Turning off stops subsequent automatic recommendation generation while preserving existing suggestions and history. Pending work rechecks the setting and permissions immediately before execution.
5. Private consultation or source recommendations are not automatically shared with other members. Changed or revoked evidence requires revalidation. Preparing a suggestion does not execute research, prove success or constitute human review.

## 34. Compact cards and work search

1. Open titles in the Projects and Research spaces sections of a grant or project. The directory uses the same cards without repeating a parent already shown.
2. Work and result counts use currently accessible entries. A work's conversation, report and execution history are not counted as separate work. A + marks a bounded count.
3. Choose List or Map and type to search Research work. Search waits for composition to finish; Enter searches immediately. Use the compact sort menu for recent activity, creation, status and other criteria.
4. Searching and changing views preserve the draft and selected work. Renaming work or adjusting map relations retains the same work ID and history.

## 35. Research kinds, requests and results

1. A research kind describes the activity: planning, literature review, model design, experiment design, code implementation, training, validation, reproducibility, analysis, discussion, writing, reporting or scheduling. Reports, code files and literature tables are result formats.
2. AI proposes a kind during refinement. Direct requests receive an initial classification; use Change research kind in the work menu to correct it. Manual choices are never overwritten automatically.
3. Every work uses Request, Results and Continue this work, in that order. Expand long requests and open or download results. Find initial requests, conversation, execution and earlier versions in History under the work menu.
4. Changing kind preserves sources, results and permissions; it does not authorize server execution or GitHub writes. Distinguish a training plan from actual training and a validation plan from completed validation.

## 36. Continue earlier goals and literature

1. Earlier paper-goal settings and discovered literature can be organized as a separate preparation work. Read its research objective and conditions as the request, and open the literature bundle as its result.
2. The literature result includes search and observation history, selection state and sources. Open or download it, distinguishing reorganized existing records from new AI analysis.
3. Earlier strategy and detailed-report IDs, conversations, reports, citations and downloads remain intact. Linking preparation as the logical starting point never invents historical runs or conversations.
4. Earlier settings and literature URLs lead to the preparation work. Its automatic execution settings retain the earlier policy; reorganization does not resume stopped AI or incur additional AI charges.
5. Renaming work never creates another map node for the same work. Work and result artifacts remain distinct; expand result versions when needed.

# 12. Research workflow, evidence and reproducibility readiness

Evidence, protocols, reproducibility records and report tools.

Open this guide online

## 1. Continue from a goal

Open Research workflow and evidence on a goal or strategy report. Inspect literature, authorized reports and code results, then continue into coding or consultation. Preview and save private hypotheses, protocols, decisions or manuscript drafts as new records. Original sources remain unchanged; changed evidence is flagged for review.

## 2. Reproducibility results and readiness

Import unchanged ReproKit v1 events using a linked project UUID as project\_key and a producer instance UUID. Retransmissions are deduplicated. Completion and numeric matches remain producer reports, not independent verification or baseline approval. Remote experiment dispatch is not connected. Record dataset, split, environment, baseline, metric, tolerance, seed, budget and access requirements, then download the evidence manifest.

## 3. Sharing, Notion and printing

Use the top share icon for team sharing or Notion publication. Find recipients by name and use ? for field help. For Notion, select an allowed parent and review the body before queuing publication. Configure a connection if none is available. Notion audiences are independent. The export icon provides full-body print, PDF and Markdown. An uncertain publication is reconciled, never blindly reposted.

## 4. ChatGPT and AI consultation

Connected tools support reading a goal workflow and recording notes or imported results. Use the lifecycle envelope inside consultant\_thread\_prepare save, followed by reviewed commit for changes. Only transmitted content is saved. Reads and record writes do not invoke extra AI or remote experiments; new consultation follows existing usage policy.

# 13. Research space guide

Getting started, four tabs, source suggestions, files, follow-ups and 32 copyable requests.

Open this guide online

## 1. Get started

1. Open Research Consultant → Research spaces. For a new space, choose a lab and research question; link a parent grant or project when needed.
2. Open new research from Overview and write one or two sentences. A title, type and status are not prerequisites.
3. Use Add materials to find evidence or attach files. Review the request, then choose Send to AI.
4. For ongoing research, open its entry in Research work and continue in the same composer. Reading and selecting evidence uses no AI; sending follows current usage policy.

## 2. Repeated work and roles

1. Continue or derive discussions, literature studies, experiments and manuscripts as needed. Keep follow-up questions for the same purpose inside that work; link evidence when creating different work.
2. New questions and notes start privately. Linking to team work does not automatically publish a private conversation.
3. Authorship and access are separate, and reviewing one version does not approve later revisions.

## 3. Existing services

1. Materials and results shows authorized parent project and grant materials. Upload files, public links or connected cloud files to the existing parent scope; the global Materials menu exposes the same record ID.
2. Spaces do not duplicate files into a separate store. External connections reuse existing connection and authentication screens.
3. Members must also have parent-scope access; space membership does not expand source permissions. Manuscripts retain their editing origin and version.
4. Live Overleaf, DOCX and HWPX editing synchronization requires separate adapters.

## 4. Continue across AI surfaces

1. Use  beside the send button in either a new or follow-up request to choose: Answer in this work, Continue in AI chat, or Continue in ChatGPT.

2. Selecting a destination runs no AI. Answer in this work sends the request. Continue in AI chat prepares your final question, selected evidence and files in the consultation panel; review and send there to start.
3. Continue in ChatGPT prepares handoff text. Copy it into ChatGPT with LUNA connected. Copying alone neither transfers nor saves anything externally.
4. Reuse the same work and conversation for the same purpose. New work is created on actual send; switching a prepared request's destination must not run it twice.
5. Only delivered and saved content and files are preserved. Check the receipt's work ID, files and download links. Choosing a destination does not broaden private conversation access or sharing.

## 5. Research decisions and reproducibility

1. Record multiple environments, datasets, metrics, baselines and repeat conditions. Existing workflow and reproducibility import uses the approved ReproKit result contract.
2. Producer-reported results, researcher decisions and independent verification remain distinct. This space does not automatically dispatch remote experiments or paid AI.

## 6. Use the four tabs

1. Overview: review work to continue, recent outputs, upcoming events and suggested research. Research work: choose work from the list or map and continue its questions, conversations and results in one view.
2. Materials & results: search sources and saved outputs and pin important materials. Open AI source suggestions here too.
3. Schedule & tasks: create or link events, GitHub tasks and report requests. Open Research team and space settings from the space menu.

## 7. Core materials and AI drafts

1. Search and pin important sources in Materials & results. Pinning and unpinning do not copy, delete or change source access.
2. Private pins enter only your own AI context; team sources still require current source permissions. Attach relevant evidence and choose Ask AI to save a request on the same work.
3. Choose Save note to preserve a private memo without AI; it does not replace the current result. Waiting, running and failure reflect actual request state; saved stages remain preserved.
4. AI requests follow administrator policy and personal limits. Reading, saving notes and pinning make no new generation call.
5. Pins and AI drafts are not factual verification or researcher approval.

## 8. Members and manuscript authors

1. Edit a member inline to load their current role and access. Only researchers with current lab and parent access can join.
2. Permission changes retain exact previews. In a manuscript record, Authorship for this manuscript stores multiple roles and order separately.
3. Authorship is versioned without changing space access. Existing space roles are not automatically copied to every manuscript.

## 9. Search, pins and space-specific schedules

1. Use the single search row and expand Advanced filters only when needed. Search and page through all accessible parent materials in Materials & results, then pin them there.
2. The pin button shows saving; the row and pin count update together after success. Use the active pin to remove it in place.
3. On failure, review the current state and error before retrying. Pinned materials shows the full list independently of the bounded AI evidence bundle, without deleting files or history.
4. Schedule & tasks defaults to this research space. Search Link existing event or task and save a private or team reference.
5. Select the full parent scope in Advanced filters. Calendar and task links retain the space filter.
6. + Event and + GitHub task open the standard editors. Request, review and apply an AI draft there; edits made after a request are preserved.
7. The existing save/send confirmation creates the real item and retains its research-space reference. Reads, links and pins incur no AI usage; generation follows administrator policy.
8. The More menu guide links directly to each step. Web, personal AI and ChatGPT use the same space IDs, references, versions and permissions; this does not imply copying an entire external conversation.

## 10. Find earlier analyses and results

1. Existing goals, code, reports and discussions retain their original IDs and permissions while being linked into one work list. Open work or search Materials & results for originals and downloads.
2. The current revision and latest saved checkpoint represent a conversation; earlier versions remain in detailed history. Distinct reports are not merged by matching titles.
3. If evidence is inaccessible, the owner may see only a preserved-result notice. Restore current connection and permissions before opening content or downloads.

4. Never infer or claim to have read blocked content. A classification-only change with identical content, version, connection and current access does not require regenerating saved results.

## 11. One work list and research map

1. Choose List or Research map in Research work. Selecting work opens the right panel. Use ↗ or Ctrl/⌘-click to open the same work on its own page with a stable URL. The default map groups results inside their owning work; Expand outputs reveals exact report and code-result cards.
2. Renaming work changes its display name without rewriting source titles, content, versions or evidence, or duplicating work. Distinct works are never merged by title.
3. Solid workflow/derivation links and dashed reference/validation links are both visible by default. Multiple origins, references and sections retain their exact report IDs and versions even when output cards are folded.
4. Counts distinguish the works, outputs and materials currently shown, not the total saved collection. Use the next work range or a focused neighborhood to navigate.
5. Renaming or reordering the list does not reposition the same works and relationships. Goal settings, literature, strategy analysis and specialized code tools remain available from that work's settings or detail screen.
6. Source changes require deliberate evidence review and refresh. Disconnecting does not delete the source or earlier history.
7. Maps, search and opening a work URL call no AI and create no work or conversation. The panel and separate page use the same work, current permissions, conversation, results, files and history.
8. While drafting, use ↗ to open a new tab. Unsent text and selected files remain in the original tab. Open in research space returns to that work. Use the original tab to resume its search, map position and draft.

## 12. Trash and a new research cycle

1. Review the target and effects before moving work or a space to Trash. Saved conversations, outputs, attachments and external originals remain preserved.
2. A new cycle archives earlier work while retaining materials and members. Restore from Trash or previous cycles without automatically restoring work individually trashed earlier.
3. Related local pending or running requests stop, but a call already sent to an external provider cannot be guaranteed recalled. Restoration itself queues no AI.
4. Original automation policies remain preserved, so an enabled policy may execute at its next scheduled time after restoration.

### 13. Navigate quickly and inspect history

1. Switch between List and Map and sort by recent activity, creation, execution state, attention or title. Changing work loads its detail in place and retains the draft question for each work.
2. Open Records and filter All, Conversation, Results and versions, Executions or Links. An answer and its result are grouped.
3. A conversation started from an earlier report or selected passage retains its original thread. Load earlier records on demand.
4. No execution requested, generating and needs attention describe execution separately from saved outputs. Downloads require actually preserved files.

### 14. Save naturally from a conversation

1. Ask the AI consultation or ChatGPT: "Save this conversation and its results to the current research work." LUNA prioritizes an explicit work, the bound conversation and a previous save receipt. Choose by title when several targets match, or preserve privately in the personal inbox before selecting a destination.
2. New work requires an explicit request; retrying the same request does not duplicate it. Review the destination, content and files before saving.
3. The personal inbox supports downloads and linking to existing work. Saving text makes no additional platform model call.
4. Files can be uploaded in resumable chunks up to 8 MB per bundle; larger materials may be saved as accessible source references. A source link is not a preserved binary copy.
5. Check the completed save receipt: uploading means only received parts are preserved. External messages and files that were never delivered are not automatically collected.
6. Team sharing and external publication retain separate scope reviews.

### 15. Full-width map and images

1. The map uses the full work area. Derivation and dependency relations guide its left-to-right layout, including independent work and loops.
2. Select a node for detail and return to the same map position. More offers PNG, JPEG and SVG exports of the currently loaded range, not a claim to cover every record.
3. No AI call is made.

### 16. Attach files and use code results

1. Use Add materials for files, images and platform evidence; paste links into your question. Upload up to 50 files totaling 100 MB through resumable chunks.

2. Originals are preserved while AI excerpts are prepared in the background. Preparation is not a full read of every original; extraction failures and text or image limits are shown separately.
3. Actually completed ChatGPT file transfers can be attached with the same permission checks. Missing code output has no download; execution history shows the failure.
4. Import preserved text code or a validated ZIP as a code changeset and inspect each diff. ZIPs are never executed or automatically installed.
5. Review repository, branch, baseline commit and changes before approving GitHub publication. Reconcile uncertain responses to avoid duplicate posts.

## 17. Keep existing work in My saved items

1. Star existing work, conversations or results to keep a shortcut in My saved items. This differs from a preserved copy of delivered conversation content and files.
2. Shortcuts reopen the same work and version; source permission revocation blocks content access. Web, AI consultation and ChatGPT use the same bookmarks and save receipts.

## 18. One library and pinned materials

1. Use All, Favorites and Saved conversations and files in one personal library. Old favorite URLs and folders remain valid.
2. Unstarring does not delete preserved originals. Materials and results initially shows pinned items; search includes all accessible materials.
3. Toggle the same pin to unpin and clear search to return to pins. Pinning does not expand source permissions.

## 19. Private work and team discussion

1. Click Only me on a work item to manage sharing. Review the exact body and selected files before sharing.
2. Earlier private discussions and personal memory are not shared automatically. Continue in team view to accumulate attributed discussion under the same canonical research work.
3. Sharing withdrawal, membership and source access are rechecked; private originals remain preserved. ChatGPT uses the same work ID and current permissions and reflects only delivered content.

## 20. Add schedules and export maps

1. Use Add to choose an event, GitHub task or report request and then create or link existing. Creation opens the standard form with optional AI drafting.

2. Linking does not duplicate the original or create an external item. Select one image format before downloading a map.
3. Dashed reference and validation relations identify evidence or verification links, separately from derivation order.

## 21. Report links and research team settings

1. Use the work or report share icon to copy its link, copy its title and link, or share to another app. A work link opens the same work on its own page; use the exact result or version link when referring to a specific output. Copying uses no AI, grants no access and still requires the recipient's current permission.
2. Review recipients, permission changes and Notion publication in sharing settings. In Research team, choose Add researchers, search names and select several candidates.
3. Review the access level and add them to the list. Authorship is optional and permits multiple roles.
4. Edit roles and access from the member row. Granting management or ending participation requires an impact confirmation; the owner retains management.
5. Basic settings holds the name, description and parent scope. If candidates are missing, check active lab membership and parent grant or project participation.
6. On a save conflict, retain your input and refresh membership. Private conversations, memory and source permissions are not automatically shared.

## 22. Work menu and next research plans

1. Open the work menu to rename a work item, manage multiple references, view history or change relevant settings. Display names do not alter original titles or result versions.
2. Find conversations, executions and revisions in one panel. Save a report next action as a research plan; private research replies can also suggest a plan within the same AI call.
3. Review the rationale, expected output, completion and stop criteria before adopting or deferring. Adoption prepares follow-up work; it does not execute AI or verify a hypothesis.
4. Continue plan with AI explicitly starts a request under current policy and allowance. Changed evidence requires review.
5. Report Next actions starts with three suggestions. Use View all for more and chapter links to inspect their evidence. Grouped suggestions retain original candidates and sources.
6. Distinguish started, deferred or dismissed suggestions and genuinely different experiment conditions. Print view and PDF omit the interactive Next actions section while preserving the report body and citations.

## 23. ChatGPT files and accurate save coverage

1. Ask to save the delivered discussion and generated files. Reuse the linked work, or keep an unresolved destination in your private inbox.
2. Only content and bytes actually delivered are preserved. Up to 50 originals totaling 100 MB use resumable upload; already preserved files are reused by ID.
3. Preserved does not mean fully read by AI. Check downloads and preparation states in the receipt.
4. Broader sharing, publication and GPU execution are not implied by a save request.

## 24. Refine a request through conversation

1. Write a few sentences in New research or Continue this work, then select sources and files in Add materials. Put links in the text. Both composers use the same refinement, attachment and destination controls.
2. Refine request opens a conversation on the right. It checks the current request, work and selected evidence, suggests a research kind and identifies missing conditions. It reuses known conditions and asks one or two important questions at a time.
3. A choice fills the reply box. Add or edit details, then choose Send reply. You may leave a condition undecided or ask for a recommendation. Repository candidates are authorized connection metadata, not proof of write permission or server availability.
4. Preview the prepared request and conditions, distinguishing your input, existing evidence, AI suggestions and unknowns. The composer changes only when you choose Apply to request. Undo restores your earlier input.
5. Applying does not execute research. Edit the result if needed, choose this work, AI consultation or ChatGPT from the destination menu, then send. You may also skip refinement and send directly.
6. Closing and reopening preserves the saved conversation. On mobile, closing returns to the original composer. An older response never automatically overwrites text you changed.
7. Starting refinement and sending each additional reply request one personal AI attempt each. Page reads, saved conversation reads and applying use no AI. Closing the panel does not cancel an already accepted AI job.
8. Administrator policy, remaining allowance, provider and time limits apply. Failures retain input and conversation without automatic paid retries. Up to 50 files totaling 100 MB are supported; uploading is distinct from analyzing all original content.
9. Selecting a saved suggestion only fills the request and evidence; Save memo stores a record without AI. Team requests use only currently authorized shared evidence.

## 25. Continue request drafts and results in ChatGPT

1. You can hand off a rough request without finishing refinement on the website. Choose Continue in ChatGPT from the destination menu for either a new request or a follow-up.
2. The handoff includes the latest composer text, selected sources, prepared attachments and existing work, conversation and result versions. An applied refinement also carries the same request session and accepted revision. Preparing and copying the handoff calls no AI, creates no work and transfers nothing externally.
3. Paste the text into ChatGPT with LUNA connected and check that the host exposes the needed tools. Ask: "First clarify the request by asking for missing conditions; proceed after we settle it."
4. Preserve the same request session, work and evidence versions when saving a request refined in ChatGPT. Saving externally prepared requests and results does not rerun platform AI. Draft storage, applying a request and executing research remain separate actions.
5. Ask: "Save the discussion delivered so far and the generated files to this work, then show download links." The platform preserves only conversation content and files actually delivered through its tools, not the entire external chat automatically.
6. A filename or temporary link is not proof of preservation. Check the receipt destination, file list, downloads and readiness. Use web upload if the host cannot deliver files.
7. Confirm a destination for new work and retain the existing work, conversation and exact result version for follow-ups. Continue through the stable work link; include the exact result or version link when reviewing a specific output. Opening a link does not save drafts, external conversations or files. Private conversations and evidence do not become team-shared automatically. Publication and paid GPU execution have separate scopes.

## 26. AI source suggestions

1. Open AI source suggestions in Materials & results. It works without pins and first reads saved suggestions without AI.
2. For fresh suggestions, check the source scope and expected usage, then choose Generate suggestions · 1 AI attempt or Suggest again · 1 AI attempt. Review titles, reasons, sources and whether full text or only summaries were checked.
3. Pin useful materials. Research with selected materials opens the composer with selected evidence and an objective.
4. Actual AI research starts when you review and send the question.

## Start with a short question

First question

Explain the attached paper's main question, method and limitations in plain language.

Select the paper in Add materials before sending.

Clarify an idea

I want to study how to reduce bias in recommender systems. Suggest three questions to narrow the idea.

Leave unknown data and equipment as questions.

Refine wording only

Refine "I want to improve a model with little data" into a request with a clear objective and expected output.

Refine through conversation, then Apply to request. Starting and each additional reply use one AI attempt; research execution is separate.

Save a note

Save this as a private note in the current work: first check uncertainty caused by the small sample.

Saving a note makes no additional platform AI call.

One next step

Identify the most important uncertainty in this work and recommend one next step with a reason.

Adopting a suggestion and executing it are separate.

Clarify a paper goal together

I want to prepare a paper. First check this space's objective and selected literature. Reuse known research questions, venue, deadline, time and resources; ask only one or two missing questions at a time. Use my answers to prepare a final request covering the literature scope and desired outputs. Do not start research or execution before I settle it.

Use Apply to request to put the accepted draft in the composer. Refinement and actual research execution remain separate.

Prepare a repository-aware implementation request

I want to apply this method to the project code. Show authorized connected repository candidates and help choose a repository and base branch. Ask only for necessary environment, implementation scope, small-example checks and completion criteria. Distinguish actually reading a repository from checking connection metadata. Do not commit yet.

A repository connection does not prove write permission or server availability. Review actual commits separately.

## Find literature and core sources

Find materials

Find materials about contrastive learning in this space and show their titles and sources.

Searching saved materials is a read within current permissions, with no new AI generation.

Recommend core sources

Recommend core sources for this research. Show saved suggestions first, and tell me the AI usage before generating new ones.

Reading saved source suggestions uses no AI; generating suggestions uses one AI attempt.

Pin selected sources

Pin the first and third materials I just reviewed as private core sources in this space.

Check the selected sources and audience; source permissions stay in effect.

Compare papers

Compare the questions, data, methods, metrics and limitations of the three selected papers. Distinguish full text you read from summaries.

A preserved file is not proof that its entire text was read.

Start from selected sources

Use the selected sources to identify unresolved questions and propose the most feasible research direction.

Use Research with selected materials, review the evidence and objective, then send.

Find references across sources

Search this space's literature collection for contrastive learning papers. Include collected and pinned Zotero references, showing each paper once. Do not start a new search run or AI summary.

Read the saved collection under current access permissions.

Read the original abstract

Show the selected paper's original abstract, DOI, authors and publication date. Keep AI summaries separate from the original abstract.

If no abstract is available, say so rather than inventing content from the title.

Explicitly request a saved summary

Request a saved LUNA AI summary and separate relevance to this space's purpose and questions, using this paper's accessible abstract or preserved body. First check source coverage and AI usage, then show the summary when it actually finishes.

A host answer and a paid platform summary request are separate. Queued does not mean completed.

Correct bibliography

I want to correct the selected reference's author name and year. Show the current bibliography first. Apply only the values I supply to this space without modifying the Zotero original.

Review field-level changes and the current revision. Recollection preserves corrections.

Export selected references as BibTeX

Export only the three selected papers as BibTeX, including authors, year, DOI and original abstracts. Do not substitute AI summaries for abstracts. Tell me if you cannot create a real downloadable file.

Claim a download only when the host actually creates a file.

Export all search matches

Export all "graph learning" matches in this space's literature as RIS across every page and report the count. Do not upload anything to Zotero.

Export is a standard bibliographic snapshot, separate from Zotero synchronization.

Restore the source bibliography

Restore only the selected paper's publication date to the source value. Preserve other edits and its work references.

A null field removes that override and restores the source value.

Use literature as implementation evidence

Link the two selected papers as references for a new model-design task. Check readable source coverage first, distinguishing applicable method details from unknowns.

Retain paper IDs, versions and permissions, and state what source text was read.

Register references as materials

Register these DOIs and paper links as references in this research space. Reuse existing materials and check the destination project/grant and audience. Do not generate new AI summaries.

Bibliography registration, original storage and AI analysis are separate; registration requires actual host tools.

Register only missing papers

Check each paper from my earlier registration request and register only items not yet accepted. Do not recreate queued materials. Distinguish storage failures from missing PDFs.

Reuse existing request receipts and material IDs.

Attach a PDF to an existing reference

Attach this published PDF to the already registered paper without creating another reference; preserve any earlier preprint. Tell me when it is stored and downloadable, and report text extraction and AI analysis separately.

Verify completion from delivered bytes and the platform preservation receipt.

Find bibliography in Materials

Find the references collected in my research space through Materials. Search by author and DOI and show each paper once after checking duplicate source identities.

Search only currently accessible references without exposing another space's private corrections.

Export corrected bibliography

Check that these bibliography corrections made through Materials also appear in the research space, then export only the selected references as BibTeX. Preserve Zotero originals and other spaces' corrections.

Read and export the same reference IDs and space corrections under current permissions.

Check original and extraction status

Check original-storage status for the registered papers. Distinguish PDFs, archived HTML, extraction failures and missing originals, and show download links for preserved files. Do not start AI analysis.

An archived webpage is not presented as an acquired PDF.

## Work through research stages

Form a hypothesis

Propose two testable hypotheses grounded in the selected literature and observations that could refute each.

Separate source claims from newly proposed interpretations.

Plan an experiment

Design an experiment for the first hypothesis, including data splits, baselines, metrics, repetitions and stop criteria.

Planning does not start an experiment or a GPU run.

Plan implementation

Inspect the linked repository's current structure, then propose small implementation steps and the necessary checks.

Label proposed code, synthetic checks and actual execution results separately.

Interpret results

Calculate changes from the baseline and uncertainty in the attached results, separating supported conclusions from claims needing more experiments.

Select the actual result file and check the calculation evidence.

Draft a manuscript section

Draft a results section using only reviewed findings in this work. Do not invent numbers or citations.

The draft does not automatically change the manuscript source or author approval.

Plan training with unknown resources

I want to prepare training for this model. First check access to the selected data, splits, baselines and metrics. I do not yet know whether a GPU server is connected and available. Do not assume missing resources exist; ask about possible environments and time/cost limits. Save only the execution plan for now; do not start training.

Unverified environments and permissions remain unknown. Saving a plan does not execute a paid GPU run.

Inspect paper-code provenance

Inspect the paper's linked code repository and report whether it is official, the inspected commit, license and environment requirements. Do not commit to my project or run it on a server yet.

Distinguish official links, inspected code, reuse rights and validation results.

Edit saved research defaults

Show this work's research conditions. Change the default evaluation metrics to the ones I confirm while preserving earlier request/result settings. Save settings only, without execution.

Check the kind-specific fields and current settings revision.

## Continue existing work

Follow up in the same work

Analyze the sample-size limitation in this work's latest result in more detail.

Open the existing work before sending to continue its conversation and results.

Continue an exact report

Use section 3 of the selected report version to prepare the next analysis plan.

Keep the selected version as evidence rather than silently substituting the latest result.

Review all suggestions

Show all saved suggestions for continuing this work so I can compare their reasons and expected outputs.

View all reads saved suggestions without a new AI call.

Derive separate work

I want to start separate research comparing another dataset while referring to this result. Preserve the link to the original work.

Keep same-purpose questions in existing work; derive work for a different objective.

Check status and partial results

Show this request's actual state, saved partial results and files available for download. Do not rerun it.

When failed or quota-waiting, inspect saved results and resume conditions first.

Append a request/result pair

Continue within the current work without creating another item. Check two assumptions in the latest result and preserve this request and answer in the same work.

Use `continuation_mode=same` and never attach an earlier answer to a pending request.

Derive from an exact result

Create separate implementation work from the selected output version. Keep the original conversation intact and link this literature and result as evidence.

`continuation_mode=new` creates the new work and originating evidence on actual save/send.

## ChatGPT, files and preservation

Hand off the final request

Check the final question and selected evidence in the copied LUNA request, then continue this work.  
Check attachment readiness first.

Verify the LUNA tools actually available in ChatGPT. Copying alone does not transfer anything.

Save discussion and result files

Save the discussion delivered so far and the actual generated files to the linked LUNA work, then show download links.

Verify actual content/file delivery and a completed receipt. Saving does not require another AI run.

No work selected yet

Save this conversation and its generated files in my private inbox. I will link them to research work later.

Actual uploads support up to 50 files totaling 100 MB. Filenames or temporary links alone do not establish preservation.

Check original downloads

List the originals preserved in this work with their readiness states and provide downloads for completed files.

Use the website upload when the host cannot deliver the file.

Refine and save from ChatGPT

Help refine the supplied LUNA request in this conversation. Preserve its request session, work and evidence versions, and ask about missing conditions. After saving the settled request, preserve the results I ask you to prepare here and the files actually generated in the same work. Do not regenerate them with platform AI. Show the saved work link, file list and download availability.

The host must expose the required LUNA tools. Check draft/result receipts and actual delivered files; do not describe this as automatic synchronization of the entire external chat.

## Team, relationships and schedules

Rename work

Rename the current work to “Small-data comparison experiment” while preserving its results and references.

Renaming does not create a new work item.

Link several references

Link the two selected works and the report as references for this work, then show their relationships on the map.

Reference, derivation and prerequisite are distinct relationships.

Share selected results

I want to share only this summary and the selected result files with the research team. Show the recipients, body and files first.

Share only the reviewed scope; private conversation history and memory are not automatically exposed.

Link an existing schedule

Link the existing next research meeting to this space without creating a duplicate event.

Reuse the existing calendar or GitHub task; review any external creation separately.

### 33. Automatic recommendations and usage

1. Open space settings in Research team to turn automatic research recommendations on or off. Existing and new spaces default to on, preparing suggestions from existing reports and consultations.
2. Opening pages, reading suggestions, searching and sorting make no new AI call. Relevant saved results refresh recommendations in the background; the first view shows saved suggestions.
3. Additional automatic AI enrichment defaults to zero attempts. Set a daily cap if needed; the space owner's allowance and administrator policy also apply. Changing settings never resets consumed usage.
4. Turning off stops subsequent automatic recommendation generation while preserving existing suggestions and history. Pending work rechecks the setting and permissions immediately before execution.
5. Private consultation or source recommendations are not automatically shared with other members. Changed or revoked evidence requires revalidation. Preparing a suggestion does not execute research, prove success or constitute human review.

### 34. Compact cards and work search

1. Open titles in the Projects and Research spaces sections of a grant or project. The directory uses the same cards without repeating a parent already shown.
2. Work and result counts use currently accessible entries. A work's conversation, report and execution history are not counted as separate work. A + marks a bounded count.
3. Choose List or Map and type to search Research work. Search waits for composition to finish; Enter searches immediately. Use the compact sort menu for recent activity, creation, status and other criteria.

4. Searching and changing views preserve the draft and selected work. Renaming work or adjusting map relations retains the same work ID and history.

### 35. Research kinds, requests and results

1. A research kind describes the activity: planning, literature review, model design, experiment design, code implementation, training, validation, reproducibility, analysis, discussion, writing, reporting or scheduling. Reports, code files and literature tables are result formats.

2. AI proposes a kind during refinement. Direct requests receive an initial classification; use Change research kind in the work menu to correct it. Manual choices are never overwritten automatically.

3. Every work uses Request, Results and Continue this work, in that order. Expand long requests and open or download results. Find initial requests, conversation, execution and earlier versions in History under the work menu.

4. Changing kind preserves sources, results and permissions; it does not authorize server execution or GitHub writes. Distinguish a training plan from actual training and a validation plan from completed validation.

### 36. Continue earlier goals and literature

1. Earlier paper-goal settings and discovered literature can be organized as a separate preparation work. Read its research objective and conditions as the request, and open the literature bundle as its result.

2. The literature result includes search and observation history, selection state and sources. Open or download it, distinguishing reorganized existing records from new AI analysis.

3. Earlier strategy and detailed-report IDs, conversations, reports, citations and downloads remain intact. Linking preparation as the logical starting point never invents historical runs or conversations.

4. Earlier settings and literature URLs lead to the preparation work. Its automatic execution settings retain the earlier policy; reorganization does not resume stopped AI or incur additional AI charges.

5. Renaming work never creates another map node for the same work. Work and result artifacts remain distinct; expand result versions when needed.

### 37. Browse the space literature collection

1. The Space literature card in Materials & results combines papers collected by literature work with pinned Zotero references. Its compact view shows three references with pinned items first.

2. View all opens a full-width literature page scoped to this research space. Each title opens its permanent detail page. Going back preserves the search context. Search by title, author or DOI and move between pages.

3. Sources with the same DOI share one collection entry while retaining original sources, versions and access. Similar titles alone never merge papers. Your entire Zotero library is not automatically published to

the team.

4. Add materials in new and follow-up requests searches these same references. Selecting one links evidence without executing AI. Unpinning does not delete the paper or its original files.

### 38. Original abstracts and on-demand AI summaries

1. Original abstract shows the provider-supplied abstract, or explicitly states that none is available. Browsing, collecting, pinning and exporting never automatically generate an AI summary.

2. Choose AI summary and relevance when needed. One request summarizes the source and separately assesses its relevance to this space's purpose and questions. One request follows current usage and access policy; queued, running and completed states remain distinct. Repeated clicks do not duplicate the same active request.

3. Use an available abstract first, otherwise an accessible preserved body or excerpt. If no text is available, add an abstract or readable source. A title is not evidence for inventing a paper summary.

4. AI summaries show actual coverage, such as abstract or body, without replacing the original abstract. AI relevance separately shows direct, partial or not-established relevance and its reason. Relevance sorting uses only current saved assessments, never inventing scores for unevaluated references. Source assessments remain separate. Source changes mark older summaries as based on an earlier source; regeneration requires another explicit request.

### 39. Edit standard bibliographic metadata

1. Use the pencil in reference details to edit title, authors, publication date, publication, DOI, original URL and abstract. Add, remove or move authors up to change order. Expand More bibliographic fields for volume, issue, pages and publisher.

2. Save applies a space-specific correction without overwriting Zotero or other external originals. Cancel returns to saved values; Restore original bibliography removes corrections. Recollection preserves manual edits.

3. Bibliography uses CSL-JSON types, creators, dates and identifiers. DOI and date formats are validated. A conflicting or invalid DOI raises an error rather than silently merging papers.

4. Current search, Add materials and new exports use edited metadata. Earlier report citations, files and source versions remain preserved. If another screen saved first, refresh the latest version before editing again.

### 40. Export selected or all references

1. Use the card download icon to choose BibTeX, RIS or CSL JSON. Export individually checked references, the entire collection, or all current search results.

2. Select all covers matching references across every page. Uncheck individual items to exclude them and verify the actual export count. References whose access has been revoked are not exported.
3. Files contain current bibliography and original abstracts. AI summaries and private research relevance never replace original abstracts. Import a downloaded BibTeX or RIS file through Zotero File → Import.
4. Export creates a point-in-time copy, not two-way Zotero synchronization or PDF transfer. Check for duplicates with existing Zotero items after import.
5. For direct delivery, open Send to Zotero, review the destination and permission to send selected citations, and apply that binding first. Select 1-20 individual references, preview the citation fields, then choose Apply reviewed changes. Existing Zotero bibliography is preserved; only new citations and collection membership are sent. PDFs, notes and deletions are excluded.
6. verified describes remote readback at the displayed verification time. For partial, uncertain or conflict results, inspect the same request instead of blindly submitting a new one. Import new items adds collected references to the space; it neither authorizes unattended AI-to-Zotero writes nor enables collection scheduling.

## 41. Continue the same work or derive new work

1. In follow-up requests, Continue in this work appends paired requests and results to the same work. New derived work creates a separate item only on actual send and retains the originating work or selected output version.
2. Report suggestions also offer Continue in this work or New derived work. Continuation and delivery destination are separate choices. Choose Save as a note in the send menu to save without running AI. Web, AI consultation and ChatGPT retain the same request ID, evidence and settings versions. Preparing, copying or changing a destination creates no work or paid execution.
3. Research flow puts earlier work below later causal work. Reference links are not silently treated as prerequisites.
4. Edit kind-specific defaults in Research conditions under the work menu. Applied refinement conditions and saved defaults are reused, while each request preserves its settled snapshot. Editing settings does not authorize GitHub writes or server execution.

## 42. Connect literature to code implementation

1. Select references in Add materials when requesting model design or implementation. AI uses the current authorized source and exact version, distinguishing abstract-only access from inspected body text.
2. Future code links retain official-versus-third-party status, supporting source, repository, commit, license, data and environment. A URL alone proves neither official authorship nor permission to reuse code.
3. The intended trace is literature → selected method → source code version → target repository/branch → patch → checks and experiments. The collection preserves evidence and extension fields for this flow.

4. This literature feature does not automatically enable training, server execution or external commits. Distinguish generated code, synthetic checks, actual experiments, independent reproduction and scientific conclusions.

### 43. Literature results, printing and collection settings

1. Literature results focus on references and expandable abstracts/reference information. Change priority or exclusion under each item's without deleting the paper. PDF links appear only when a known URL is available.

2. Print view opens every reference matching the current filter, not just the visible page. Search, settings and action controls are excluded from printing.

3. Manage recurring collection and analysis in the work's → Automation settings. Collect literature now requests one additional collection without enabling automation; current access, query conditions and queued status are rechecked.

4. Collection updates metadata and source observations. It never automatically generates per-paper AI summaries and preserves earlier collection and analysis history.

### 44. Find and edit literature in Materials

1. Open References in Materials or search by title, author or DOI. References saved in a research space are discoverable in the same catalog; do not register another copy.

2. Select a card title to open its wide permanent detail page; View all opens the dedicated literature list. Translation, sharing, favorites and print are at the top, followed by authors, venue, publication date, abstract and source links. Original abstracts and AI summaries are separate. Generate an AI summary only with the explicit button or request.

3. Edit only the required bibliography fields. Materials and the space use the same reference ID and correction. Check the applicable research space; other spaces and Zotero originals are not silently overwritten.

4. Select references and choose BibTeX, RIS or CSL JSON under Export. Distinguish selected references, all search matches and the complete current scope. All includes every page in that scope, not just the visible rows.

5. Use File → Import in Zotero to read the exported file. This is a bibliographic snapshot, not automatic two-way synchronization. Unpinning does not delete the original material.

### 45. Register references and request source storage in conversation

1. Tell native AI or ChatGPT the research space and the DOI, source URL or bibliography to register. Example: "Register these papers in this space, reuse existing references and do not generate new AI summaries."

2. Check the destination project/grant, audience, new/existing references and original attachments. The host must actually expose the required tools. Search-only tools cannot prove registration.
3. Check each returned material link and status. With a linked research work, the actual registration request and receipt also appear in its history. Ask for a destination if work context is unclear; never invent earlier conversation that was not delivered. Bibliography registration, pinning, source storage, text extraction and AI analysis are distinct. Reuse the same request ID when retrying to prevent duplicates.
4. Add a PDF later from the reference detail page. Active authorized participants can contribute files, preserving contributor and file versions. Identical files are not linked twice and earlier PDFs are not overwritten. Distinguish published papers, preprints and supplements. In conversation, identify the existing reference before attaching a PDF. Example: “Attach this PDF to the existing reference without creating another paper.” Only delivered files or accessible originals can be stored.
5. A ZIP or download link generated in ChatGPT is not proof of platform storage. Check file delivery and the platform receipt. Registration, reading, pinning and exporting do not themselves generate AI summaries.
6. The tool sequence is `consultant_material_prepare` → `consultant_material_commit` → `consultant_material_status`. Review the destination and audience, then pass the unchanged `preview_token` and review. A cached host can use the existing `consultant_thread_prepare` save / materials envelope for the same service. If neither route is exposed, refresh the connection tool list.
7. A request accepts up to 50 items and 100 MB of combined attachments. Distinguish “register bibliography only” from “also store an accessible PDF.” A paper landing page is not reported as a saved PDF; supply an actual PDF file when needed.

## 46. Understand storage, PDF and extraction states

1. Once registered, bibliography can be searched and edited before source collection finishes. Do not create duplicate references because earlier files are still processing.
2. Queued and running storage are distinct. Stored means a verified preservation receipt exists. An archived HTML page is not a PDF.
3. If the original is stored but text extraction fails, it remains downloadable. Combined source attachments may total 100 MB; safe text extraction supports up to 20 MB per file. Larger originals are preserved with an extraction-limit status. Scans, protected documents or unsupported formats may not yield readable text.
4. Retryable storage failures use bounded attempts and backoff. For persistent failures or access requirements, inspect the reason and supply an accessible file. Access and login restrictions are not bypassed.
5. A PDF downloadable from its reference is distinct from source archiving and text extraction. Linking a PDF does not prove external synchronization, extraction or AI analysis. AI analysis is a separate request

with its own usage. Stored originals or extracted text do not mean AI analysis has completed.

## Bibliography and PDF visibility

New public-paper bibliography is visible by default to active signed-in platform users, not anonymous visitors. PDF originals, private analysis and research spaces retain their existing access. Reusing an existing private source does not publish it. Edit bibliographic fields and grant/project/researcher metadata on the permanent literature page. External Zotero originals are not overwritten.

## Find and attach PDFs from ChatGPT

Example: “Find references in this space that lack a PDF. Locate public originals, verify the title and DOI, and attach them to the existing references while preserving earlier files.” Tools import a public HTTPS PDF or an actually transferable ChatGPT file and verify bytes, size and hash. They never bypass login or paid access. Original preservation, text preparation and AI analysis are separate states; an unavailable file transfer is not a saved result.

## Compare literature and save to the same work

Example: “Compare these three papers for the Paper Goal and Literature Survey work. State the PDF version and actual coverage; tabulate model structure, data, metrics and limitations. Save the delivered analysis as another request/result pair in this same work and attach the comparison file.” The existing work ID and earlier results are preserved. An excerpt is never labelled a complete-paper reading. Specify explicitly when a new derived work is intended.

## File transfer and retry

Example: “Save the actual image and BibTeX files just created as outputs of this work. Report each saved status and download link.” Attach up to 50 files, combined 100 MB. Reuse the request key for retries; reselect files to renew expired transfer links. The file tool must actually be available in the ChatGPT host. If unavailable, transfer the files through Add sources on the platform. A filename or temporary link is not a preserved file.

## Edit, export and apply requests

Search, select required references or all references, then use the download icon beside search for BibTeX, RIS or CSL JSON. Example: “Correct the selected authors and DOI while preserving the original.” AI summaries run only after a button action or explicit request. Applying the refinement preview updates both the current composer and its local draft, ready to inspect before sending. Dashed graph edges indicate references; solid edges indicate research flow.

## Start paper analysis in the same work

In the existing work composer, use Add sources to select one reference or PDF. Enter a goal-specific question, choose Paper analysis and discovery in the Send to AI menu, and send. Source reading, backward/forward/topic search, registration, synthesis and manuscript drafts are checkpointed. Search uses the public paper title or your explicit public query, not private goals. Verified references and final recommendations are registered automatically; PDF collection is queued with actual storage status. Missing originals limit analysis to the abstract. Forward and similar searches cover the first 25 results each; request further searches separately. Paywalls and institutional login are not bypassed; unavailable originals retain a reason. Registration, original collection, AI analysis and author review are separate states.

## Goals, notifications and collaborative review

Use work menu → Research conditions to edit the question, venue, deadline and constraints. Expand Collaborative guidance to opt into a reviewed goal snapshot and next-action notifications. Only the published goal version is shared; private edits do not silently replace it. Notifications open the same suggestion ID used by the composer and do not execute AI. Deferred or dismissed suggestions are not resurrected. Before joint use, check work sharing, access to originals and the shared goal separately. Approved methods also require use and provider-processing permission. Participants use the research team destination for follow-up analysis. Private results are not published automatically. A researcher confirms submission or goal completion.

## Improve v1 into v2

Use work menu → Know-how settings to select exact approved versions. The default direction and relevant common/activity guidance are pinned; four stages are not required for every work. Existing work-only know-how remains available in collapsed history. Select currently accessible, non-deleted request/result evidence and distinguish paper-specific fixes from reusable improvements. Compare against the actual active v1 using matched input/model/work conditions, held-out input and human review receipts before considering v2 promotion. Promotion changes future defaults and is reversible; it does not rewrite existing results. The built-in method before initialization is distinct from a reviewed and evaluated v1. Method versions, manuscript revisions and model IDs are distinct. Output-only use, source reading and derivation are separate permissions.

## Delete and restore request/result pairs

The author or an authorized space manager can move an accessible request/result pair to trash through its menu. Removed content is excluded from new analysis, method improvement and downloads; other researchers' follow-ups and existing quotations remain. Restore through Deleted requests and results; this never restarts AI. Management does not grant access to another person's private conversation. External copies and provider costs already incurred are unaffected.

## ChatGPT, consultation, Overleaf and costs

Continue in ChatGPT preserves the same work ID and selected evidence, saving only delivered results. Survey request/status/resume/external, method reads/improvements and pair delete/restore reuse consultant\_thread\_prepare/commit reviews. Independent host analysis never claims to have executed a protected method. Core instructions are excluded from ordinary results, history and exports. Provider processing permission is reviewed separately; this does not guarantee that outputs cannot reveal or imply knowledge. Download the draft and BibTeX, review them, then incorporate them in Overleaf; there is no automatic bidirectional sync. Per-request, daily and provider limits pause work while preserving completed stages. Inspect saved stages and usage before explicitly resuming. Team consultation tools currently support requests and status; use the web for resuming and method management. Host tool availability, new server PDF collection, AI quality and Lightsail latency require separate verification.

## Keep only the know-how you choose from a discussion

In the same work composer, ask: “Separate the literature-analysis and manuscript-writing know-how from this discussion. Propose only the changes to each existing method.” Only delivered, saved requests and results are evidence. External ChatGPT or Codex can prepare structured proposals using returned evidence IDs and current permissions, then use the existing reviewed save contract. On the proposal page, select only the drafts you want to keep and review the save. Unselected proposals and active versions remain unchanged. Preparing or saving a draft does not validate quality or approve v2.

## Apply several methods together

Use the applied know-how count below the work title, or work menu → Know-how settings, to manage defaults. No empty badge appears for zero items. Search names, handles or activities and review exact versions and roles before saving. Add to a work on the independent detail page lets you choose a space and work and pin an exact version. Type @ in a follow-up, new research request or personal AI consultation to search researchers and know-how together. Choosing know-how adds it only to that request without changing defaults. [[ source citations keep their existing behavior. Multiple selections never auto-merge assets or adopt newer versions. Web, consultation and ChatGPT handoff check the same IDs and current permissions.

## Split, merge, originals and shared evidence

One discussion can produce independent proposals that reference the same segment as shared evidence. Drafts start privately. Separate paper-specific corrections and unverified scientific conclusions from reusable guidance. Change organization prepares a split or merge request in the existing composer for review and sending. A merge proposes a new asset linked to exact parent versions. It does not delete or overwrite the versions, history or active state of assets A and B. Changing how a proposal is organized requires a newly reviewed proposal, not an unreviewed browser-side rewrite.

## Version inheritance and evaluation

Derived work inherits the exact selected versions of independent assets. For existing work-specific methods, the collapsed Work know-how controls retain version and candidate selection in the original work. Review generalized instructions and create a separate asset before reuse in other work. Execution applies modules relevant to the current stage, such as literature, synthesis, writing or review, while retaining links to other stages. New versions are not adopted automatically. Initial approval reviews the actual instructions and their generalization; it does not claim measured quality improvement. Improvement candidates require actual runs with matched questions, inputs, models and effort, held-out cases and human review before adoption. A candidate trial is a one-off selection that does not change the default and uses the normal AI allowance.

## Protection, storage and execution boundaries

An independent detail page shows introduction, ID, handle, authors, usage, current permissions and versions. Editing the introduction or image link does not change execution guidance; image access is checked separately. The default direction is protected guidance. Write a draft or propose a direction improvement, review it and select which drafts to keep. Saving an introduction or candidate does not change the active version. Ordinary catalogs and proposal cards exclude protected instructions. Source review requires current permission and AI-provider processing permission is separate. Reads, selections and structured proposal saves do not run AI, while sending analysis or reorganization requests may use your allowance. Independent external analysis does not prove protected-service execution. Instructions belong in encrypted backups, not code or static files. AI quality, actual ChatGPT tool availability and server speed require separate validation.

# 14. Operating grants and projects

Set up grants and projects and manage GitHub research work.

Open this guide online

## 1. Create a grant and its projects

1. Open Research grants → Create grant. Enter the title and description and save.
2. On its homepage choose Create project, enter a title, type and description. This project belongs to that grant.
3. If a needed type is missing, ask a platform administrator to add it in Lab management → Project types.
4. Assign participants and managers. A project can exist without a source connection; its creator becomes the responsible manager.

## 2. Read the project overview

1. Compare child-project summaries on the grant homepage, then open a project title.
2. Review recent reports, events, task status and the material library. Material cards open filtered lists.
3. Grant totals include accessible child-project materials. If a total differs from expectations, check filters, classification, permissions and collection time.
4. Manage authentication and collection under External source settings; edit research content in GitHub, Notion or Drive.

## 3. Work with GitHub tasks

1. Record the objective, owner, due date and acceptance criteria in a GitHub issue and add it to a Projects board.
2. Connect the repository and review discovered boards. For zero boards, verify the actual board URL and organization Projects read permission.
3. Review status, owners and dates. Where task editing is available, verify remote write permission and the target board before saving.
4. Collect again to verify the remote state. Done or a closed issue does not prove experimental success or PI approval.

## Change GitHub Start date and Target date

1. Open Research work on a grant or project, then choose Change work. Explicitly select Start date or Target date from the board fields.

2. The last collected date appears. Change it as needed. Saving an empty date clears that date on GitHub.
3. Confirm the remote change and select Save to GitHub. A matching date in the GitHub response produces a confirmation. The local list refreshes after collection. If only acceptance is confirmed, check Recent work change requests.
4. Concurrent changes to the same field or board configuration are not overwritten. Check GitHub and collect the latest values before trying again. Unrelated comments or other field changes do not block the date edit.
5. New work at grant level offers grant-wide and child-project boards you can manage. Enter Start date and Target date separately. Check for an already-created issue after partial failure; the platform does not automatically create duplicates.

#### 4. Create an issue from project research work

1. Open Grants → grant → project → Research work and choose New work/Create issue. You need project management permission and a write-capable GitHub connection.
2. Select the actual destination board and verify its repository.
3. Enter a title, objective and acceptance criteria; select multiple researchers and check their GitHub IDs. Each must be assignable to issues in that repository.
4. Choose Status, Priority, start/due dates and supported board fields. Start date defaults to today. Size is the board's work-size value, not headcount or automatically calculated hours. If Priority only offers Unspecified, check the board options/mapping and collect again.
5. Review related scopes, researchers and tags. Confirm the remote write and save, then inspect the actual issue, board membership, assignees and fields on GitHub.
6. For partial or uncertain results, check GitHub first. Do not create another issue without checking whether the first was created.

#### 5. Shared sources, wiki links and scoped popularity

1. If child projects share GitHub, Drive, Notion or Zotero sources, connect them once at grant level. Connect project-specific sources to that project.
2. Keep important sites and reference links with explanations in the grant wiki, including project-specific links. Adding a wiki link is separate from connecting external sources or collecting document content.
3. Grant popularity is scoped to that grant; project popularity is scoped to that project. It is different from lab-wide popularity.
4. Detailed progress/work/material summaries may load when expanded. Check search filters, membership and deletion state for missing entries.

## Report scope on grant and project pages

1. The progress summary and recent reports on grant/project pages refer to reports linked to that scope. Similar titles or the same author do not make another scope's report a report for this project.
2. Open report details and check the target scope and evidence affiliations. Distinguish a whole-grant aggregate from an individual project's report.
3. If no report belongs to the scope, check source links and generation scope. Do not substitute an unrelated report as evidence of this project's outcomes.
4. For an incorrect association, send support the page and report URLs. Check the linkage before trying to resolve it by deleting a report.

# 15. Grant wiki · organize important knowledge

Write, propose and review wiki pages; recover from storage errors.

Open this guide online

## 1. Open the wiki and choose a structure

1. Open a grant or one of its projects, then Grant wiki → Open wiki. All child projects use the same grant wiki.
2. The wiki does not automatically copy every source. Create topics people repeatedly need: research overview, methods, data and analysis, results and decisions.
3. Create an overview first. Expand Placement and edit reason and select Use as wiki start page. Keep it at the top level. Pin other important pages.
4. Current wiki editors in Editor settings lists names, usernames and access roles. Editors use Create page to write and publish directly. Other participants use Propose a new page to submit a title, content, sources and reason. Pending proposals are visible only to their author and reviewers. Approval publishes a new top-level page; rejection creates no page.

## 2. Write a page and curate key sources

1. Choose Create page and enter a title and short description. Select a parent to nest the page, for example preprocessing and evaluation under Methods.
2. Write Markdown, using the heading, bold and table tools. Check Preview. Headings from # to ##### produce a table of contents automatically.
3. Mention researchers with @username and search accessible documents with []. For a bibliography entry, link the paper title to its actual DOI address and add the author, year and explanation. Use the editor link tool and check Preview.
4. Under Collect key sources, search by title and select a result. Explain its importance in the body. You can also use [] directly, with one citation per line. Limits are 30 key sources and 12 inline citations.
5. Source cards keep links rather than copying entire documents. Inaccessible sources show an availability message without revealing their titles. Manually pasted prose is not automatically redacted: do not paste private research originals.

### 3. Images and publication

1. In a new page, Edit page or an edit proposal, place the cursor in the body and choose Add image. A JPG, PNG or WebP is inserted at that position; saving a draft first is no longer required. Files must be no larger than 1.5 MB and 16 megapixels; stored images are resized to at most 1600 pixels per dimension. SVG and automatic remote images are unsupported.
2. A page permits up to 20 new files and images in its body. Upload starts are limited to 50 per user per day and three unfinished uploads at a time. Images require page permission; ordinary readers can access only images in the current body.
3. Publish shares with currently authorized grant and child-project participants. Only current editors see drafts. Saving a published page as a draft hides it from ordinary readers.
4. Bookmark published pages and discuss them below. Body @mentions link researchers; saving the body does not send alerts. Use discussion comments and mentions when notifications are needed.

### 4. Find, collaborate, restore and export

1. Search titles or descriptions. Lists show 18 pages at a time; child-page links navigate the hierarchy. The start page and recently organized pages appear on the right.
2. Grant and lab managers can edit directly and review proposals by default. In Editor settings, search by name or @username under Additional reviewers and editors, select people and save. Remove a selected chip with × and save to revoke additional access. Participants can still propose edits. Additional editors cannot grant access to others. Losing grant access also revokes wiki access.
3. In Revision history, select a version to view its body and compare against the current page. Restore as a new draft preserves history and creates a new revision. Review and republish it. Renaming a page preserves its URL.
4. Concurrent edits show a conflict rather than overwriting a saved change. Keep your input separately and merge it with the latest version.
5. Trash preserves content and history. Move or trash child pages first. Open the Trash tab and Recover as draft.
6. Use the top share icon to share to a channel or copy a link, and the star to add or remove a favorite. Recipients still need permission to read the page. The export icon offers print view, PDF, Markdown, text, and Markdown and images ZIP. ZIP includes the current page and used images. Only currently accessible citations are exported; relative platform links refer to the original site. Wiki bodies, images, revisions and proposals are stored in the lab's preconfigured Google Drive. The server retains permissions, hierarchy, Drive references, titles and descriptions. Open Editor settings → Wiki storage in Google Drive to check pending migration counts. Existing content leaves the live DB only after upload and download verification. Keep your input and retry when a storage error appears; a failed save is never reported as successful. Edit through the platform: changing or deleting Drive files directly fails integrity checks. After server loss,

restore the operational DB and separately held recovery keys first, then use the Drive wiki index. Drive content alone does not restore authentication or permissions.

7. All grant and child-project participants can select Propose an edit on a published page. Enter the proposed title, content, sources and reason, then submit. The original remains unchanged. Edit proposals lists your own proposals, or all grant proposals for reviewers. Reviewers compare the changes, preview content, add a review note and approve or reject. If the original changed in the meantime, submit a new proposal based on its latest version. Approval adds a new revision.

## Insert files and images inside the body

1. Open a new wiki page, Edit page, a new/edit proposal, Write report, or a manual magazine article.
2. Select the Notion connection for a report or the lab for an article first. That lab needs a configured private Google Drive archive.
3. Click the insertion point inside the body. Choose Attach file for PDF, DOCX, HWP, TXT, Markdown, ZIP and other supported files, or Add image for PNG, JPEG or WebP.
4. Limits are 20MB per file, 1.5MB per image and 20 attachments per document. Once Drive storage is confirmed, a download link or image Markdown appears at the chosen position. Images are normalized to safe JPEGs.
5. Check Preview, then save/publish or submit the proposal. Use Retry attachment or Cancel attachment if transfer fails. If Google archive authorization has expired, keep the editor open, use Open archive account settings in a new tab, reconnect the existing account and retry. No new folder is needed. Ask your lab administrator to reconnect if you cannot manage settings. Finish uploads before saving.
6. Before saving, only the uploader can access new files. Pending proposal files are available only to the author and reviewers. Saved files follow the document audience, including Only me and revoked memberships.
7. Notion exports and wiki ZIPs retain authenticated platform links for new attachments. They do not make files public or copy large originals into the ZIP. Download originals using their links. Read-only collected documents are still edited in their original service.

## Reconnect an unavailable wiki store

1. If the list opens but the body says Wiki storage is unavailable, keep any unsaved text separately. The error alone does not mean the document was deleted.
2. Open Lab management → AI/Drive preservation → Connect Google Drive archive account. Sign in and consent with the account holding the existing wiki. This is separate from project Drive reading and personal file-picker authorization.
3. Do not create a new archive folder or change its address just to reconnect. Reopen the original page and check its body and revision history.

4. If it still fails, ask the administrator to check folder access with that account, preservation enablement, storage and revoked Google access. Do not overwrite an uncertain result with a blank page.
5. OAuth apps in Testing can have authorization expire after seven days. For recurring failures, ask the administrator to check publishing status and Google policy. Never include passwords or tokens in a ticket.
6. If consent returns “Another lab operation is in progress,” your existing connection was kept and the new authorization was not saved. Wait briefly, then start a new Google account connection from the same AI/Drive preservation settings. Do not refresh the failed callback or reuse its authorization URL. You do not need to stop ongoing research work.

Official guidance: <https://developers.google.com/identity/protocols/oauth2#expiration>

# 16. Collection and synchronization

Distinguish source registration, authorization and completed collection.

[Open this guide online](#)

## 1. Registration versus collection

1. Open Grants → grant or child project → External source settings. Consider a grant-level connection for shared sources.
2. Choose Add connection and a provider. Follow the provider-specific GitHub, Notion, Google Drive or Zotero chapter for authentication.
3. Saving a connection does not complete collection. Click Collect now on its card and refresh the status.
4. Check the latest time, count and status, then open a collected item and compare its title and body with the original.

## 2. Keep collecting changes

1. Open connection settings, enable the connection and periodic collection, and set an interval such as 60 minutes.
2. A lookback window can exclude older documents. Consider an unrestricted initial window for historical plans.
3. Edit the source and wait for collection or click Collect now. Collection time and source modification time are different.
4. Inspect the new version and body. Recollecting unchanged content is not a new research result. Periodic collection does not create AI report schedules.

## 3. Failed, partial or stopped collection

1. Expand Collection history and details on the connection card and read its error code.
2. For authentication\_or\_configuration, check source permissions, token expiry and Google authorization. Entering a folder URL alone is not Google authentication.
3. Partial collection can reflect format, body-size, request or descendant limits. Verify required documents rather than treating it as complete.
4. To pause, disable the connection or periodic collection and save. Deleting a connection does not delete the external repository, Notion page or Drive file.

## 4. Grant-level connections and personal accounts

1. Open a grant's External source connections. A connection manager can add a provider here for shared sources, or at project level for project-specific material.
2. Use a recognizable label such as Shared Zotero references or Project design Notion. Never put tokens or internal secrets in the label.
3. During Google authorization choose your account with access to the actual folder/calendar. Manage separate accounts through separate cards rather than borrowing another researcher's authorization.
4. GitHub organization administration does not replace repository/Projects token permissions, organization approval or SSO authorization. Save, collect and open the resulting material.
5. For a shared GitHub connection choose GitHub → Repository or Projects board. Enter <https://github.com/OWNER/REPO> for a repository, or <https://github.com/orgs/ORG/projects/NUMBER> and the repository to collect for a board. Personal boards use a users URL. Use a recognizable label and put the token only in the token field.
6. Save → Collect now → the grant research work page to inspect boards and material. Child projects also show the shared connection without copying it. Grants without child projects work too. Grant owners, delegated grant managers and owners or delegated managers of active child projects can add, edit, authorize and collect shared connections. Ordinary members cannot manage connections. Delegation is an actual management permission, not a role label. Grant settings, membership administration and common board field/work changes still require grant management rights. Remote writes also require token permissions and explicit confirmation.

## 5. Overleaf and OneDrive source links

1. Choose Overleaf or OneDrive under Add external connection and enter a recognizable label and source URL.
2. Save and open the link in the external service. Viewers also need access there.
3. Direct connections currently store source links. They do not automatically collect every file/PDF or turn them into AI input.
4. To read an Overleaf GitHub copy, arrange authorized synchronization in the source service, then connect and collect the GitHub repository separately. Check mirror freshness at the source.

## Root-page and nested child documents

Direct child pages on the connected root and nested child documents are collected too. Non-database child documents appear as reports by default; database entries retain their configured database roles. The root itself remains an overview material. Reports published by the platform retain their existing report IDs without duplicate registration. Automatic collection preserves manual classification. An explicitly requested administrator repair compares the current preview and changes classification while preserving notes, tags, researchers, dates, original text and history. Source permissions and connection scope do not expand. Large trees continue across batches: distinguish the last check from a completed scan. No additional AI call is made.

## Resume collection and map databases

Notion and OneNote save unfinished scans and continue them later. Last check and last complete scan are separate. Open Notion database classification and coverage, select roles for discovered database IDs and preview changes to existing materials. Manual edits and material IDs remain intact. Denied source access restricts reading; a single missing inventory entry does not delete history. Optional attachment preservation uses the existing private archive within size limits and makes no AI or OCR calls.

## Connect a Slack channel

Choose Slack and a channel URL in the grant or project source settings, then save. Ask a workspace administrator to approve the internal app and invite it to that channel before connecting Slack. Missing server app settings remain pending. Only messages and threads in the selected public or private channel become materials. DMs, message sending and automatic AI reports are excluded. Attachments retain source references and file metadata; this does not mean their bytes have been preserved. Collection makes no AI calls. Slack plan, retention and API limits still apply.

## Connection recovery and progress

Full traversal, text collection and attachment preservation are separate states. Rate limits and temporary outages resume at the displayed retry time, with longer intervals after repeated failures. Reconnect the affected account when consent has expired or been revoked. Google source reading and Drive preservation have separate grants. Google Testing status or administrator restrictions cannot be repaired by server retries alone. Check the connection code and app administrator settings. Collection runs separately from AI work and never enables disabled automation.

# 17. Connect a GitHub repository

Connect a repository and verify collected commits and documents.

Open this guide online

## 1. Find the repository URL

Sign in to GitHub → organization or personal profile → Repositories → open the repository → Code. Copy `https://github.com/OWNER/REPOSITORY` from the address bar, without `/tree/main`, `/issues` or `/projects`.

Check the branch selector above the file list (for example `main` or `master`).

In LUNA choose GitHub → Repository. Enter a descriptive connection name, paste the repository URL, enter that branch and use `.` for the entire repository folder.

## 2. Create a read token

GitHub avatar → Settings → Developer settings → Personal access tokens → Fine-grained tokens → Generate new token.

Set Token name and Expiration. Resource owner must be the owner organization/account. Under Repository access choose Only select repositories and select the target.

Under Repository permissions set Contents, Issues and Pull requests to Read-only; Metadata is automatic. Click Generate token and paste the result only into LUNA's read-token field. An organization approval request must be approved by its administrator.

Tokens also avoid the lower unauthenticated rate limit for public repositories. Do not paste a token into chat or documents.

## 3. Save, collect and verify

Enable the connection, optionally enable scheduled collection and set an interval such as 60 minutes. Save connection → Collect now → Refresh status.

Open the project home → material library → Source code or Materials, then inspect a file's source link and update time. Partial collection means supported scope or request/body limits applied; open the collection details. Add another connection for each additional repository.

# 18. Connect GitHub Projects and manage work

Connect a Projects board, map fields and create work.

Open this guide online

## 1. Create and identify the board

GitHub organization → Projects → New project → Table or Board → name → Create project. For an existing board, click its name in Projects.

Copy the opened board URL <https://github.com/orgs/ORG/projects/NUMBER>. A repository /projects URL is only a directory, not a board.

In LUNA open Research work → Linked boards → Add GitHub Projects. GitHub / Projects board is preselected. Enter a name and board URL; Repository to collect uses OWNER/REPOSITORY. Each grant or project uses one board. Edit or disable the previous connection before replacing it.

## 2. Board and repository permissions

GitHub avatar → Settings → Developer settings → Personal access tokens. For an organization fine-grained token, select its repository, grant Contents read and Issues read/write if creating or changing issues. Under Organization permissions grant Projects read or read/write for changes. Complete organization approval when required.

If fine-grained tokens are unsupported for your account/board operation, follow the official Projects API instructions: classic read:project for reads or project for writes, plus repo for private repositories. Authorize Configure SSO when required by the organization.

Paste the token into the connection token field. Explicit create/edit actions also use this token, so writing requires the corresponding write permissions.

## 3. Map fields and create work

Save → Collect now → Research work → Linked boards. Open Field mapping and map Status, Due date and Priority to their corresponding platform roles. This changes display/aggregation interpretation only.

To add a missing date field on GitHub: Table → + at the right → New field → Due date → Field type Date → Save. Collect again in LUNA to refresh the field choices.

Create work → title and body → Status such as Todo/In progress → due date, priority and assignees → confirm GitHub write → Save. LUNA creates an Issue, adds it to the board and applies fields. Assignees need a profile GitHub ID and must be assignable on GitHub.

If the result needs review, open the created GitHub issue before retrying. Uncertain create operations are not retried automatically.

# 19. Connect Notion

Authorize a root page and synchronize Notion documents.

Open this guide online

## 1. Create an integration

Sign in as a workspace owner at <https://www.notion.so/profile/integrations> (Creator dashboard). Build → Internal connections → Create a new connection → enter a name and select the workspace → create. Older interfaces may say New integration.

In Capabilities enable Read content. Insert/Update content is unnecessary for collection alone. For LUNA publication also enable Insert content and allow access to the publication parent. Copy the Installation access token (formerly Internal Integration Secret) from Configuration and paste it only into LUNA's token field.

## 2. Grant access to the root page

Open the top-level project page in Notion. Top-right • • • → Connections → Add connections → search for the integration you created → confirm. Menu wording can vary by Notion language/version.

Share → Copy link. In LUNA select Notion and enter the connection name, page URL and token. Registering a URL alone does not grant Notion access.

Child pages are traversed where permissions inherit. Add the connection on separately restricted child pages too. For another workspace, create an additional LUNA connection using that workspace's integration/token.

## 3. Verify new and updated documents

Enable the connection and optional scheduled collection → Save → Collect now. Search Materials by this project and the Notion page title, then inspect its body.

For a test, add a short sentence to the connected Notion page and collect again. Verify the source update time and new content; remove the test sentence in Notion yourself.

For zero results/access errors check the token workspace, the page's Connections menu and private child-page access. Large trees may require further collections; unsupported blocks and limits appear in collection details.

## 20. Connect OneNote research notes

Authorize Microsoft read access and synchronize a selected research section.

Open this guide online

### 1. Configure the Microsoft app

The platform administrator configures the registered Microsoft Entra application. Register the exact OneNote callback URL shown in LUNA as a Web redirect URI; it is separate from the OneDrive callback. LUNA requests delegated Microsoft Graph Notes.Read, offline\_access and openid, not app-only authentication or Notes.ReadWrite. Organization policy may require administrator approval. App registration alone does not connect an account. Enter secrets only in platform settings, never in chat or documents.

### 2. Consent and choose a research section

Project or grant → External sources → OneNote → Save → Connect Microsoft account. Sign in with your own Microsoft account, inspect notebook and section names, and select one research section intended for this space's existing audience and archive policy. Microsoft consent is broader than one section; LUNA limits collection to your selection. Do not import an entire private notebook into a project or shared lab archive. A new sign-in requires selecting the section again.

### 3. Collect and verify

Choose Collect now, then compare the title, source update time and body against the original. Text and tables are read; images and attachments have separate reference entries. Supported asset bytes are privately preserved only when original-attachment preservation is enabled for this connection, within bounded size and current permissions. Handwriting, equations, free-form layout and OCR are not guaranteed; generative AI is not called automatically. Partial collection or an unpreserved asset is not full completion.

### 4. Scheduling, retries and disconnect

The server worker performs scheduled reads; your PC need not remain on. Microsoft's per-app/per-account hourly limits may require several batches for the first import. Check continuation, next retry time and the last complete inventory. Repeated errors stop automatic retries; resolve the cause and choose Collect now. Revoked consent or changed app/account authority requires reconnection. Disconnecting in LUNA stops collection and access using that credential without deleting OneNote originals or changing Microsoft account permissions.

## 21. Connect Google Drive

Select a Google account and check folder collection and extraction.

Open this guide online

### 1. Choose a folder

Sign in at <https://drive.google.com> with the account you want. New → New folder → name it after the project, or open an existing folder. Upload or move your documents there.

Copy [https://drive.google.com/drive/folders/FOLDER\\_ID](https://drive.google.com/drive/folders/FOLDER_ID) from the address bar. To connect one file, use its Share → Copy link.

In LUNA select Google Drive, a descriptive name (including the account), paste the URL and choose Folder/subfolders or file. Whole Drive imports accessible personal and shared files into this project; choose it only when that entire scope is intended.

### 2. Select and authorize the account

Save connection, then click Google account connection on its card. In the Google account chooser, select an account that can actually open the folder. Use another account allows a different Gmail/Workspace sign-in.

Review and allow Drive read permission, then return to LUNA. If the OAuth app is in testing, ask the administrator to add that email as a test user. `access_denied` may indicate institution policy/test-user restrictions; `redirect_uri_mismatch` requires the administrator to correct the OAuth callback.

For different Google accounts, create separate connection cards and authorize each with its intended account. Reauthorizing an existing card with a different account replaces that card's credential.

### 3. Verify documents and synchronization

After authorization choose Collect now. Supported text extraction includes DOCX, HWPX, PDF, TXT, MD, PPTX, XLSX and Google Docs/Sheets/Slides. OCR of scanned PDFs and video transcription are separate from this document reader.

Materials → filter by project → open a file → inspect body/source update time. Scheduled collection revisits new and changed files.

For Partial or Continue scan pending, inspect details and wait for subsequent collections. Original Drive files are not modified or deleted.

### 4. Analyze PDF and image content

1. Open the collected Drive PDF/image. Ordinary text extraction may miss scans, diagrams and tables.

2. Automatic Drive analysis is enabled by default and processes existing and new files sequentially. Manual requests are available under File visual analysis. Source access and the connection owner's AI permission and daily allowance apply. See Personal AI memory and source coverage for supported formats and limits.
3. Compare completed notes with the original, especially figures, tables and numbers, before using them in research decisions.
4. Automatic analysis runs sequentially; it does not finish every file immediately. Format, size, page limits, AI allowance and access can restrict analysis. Check freshness after the source changes.

## 22. Connect Zotero references

Import references from a personal or group library.

Open this guide online

### 1. Identify library and collection

In Zotero desktop click Sync at the upper right to update the web library. Sign in at <https://www.zotero.org> and open the personal or group Web Library.

For a group open Groups → the group: the number in <https://www.zotero.org/groups/NUMBER/NAME> is its library ID. For a personal library use userID shown under Settings → Feeds/API → API Keys. Do not enter the username instead of its numeric ID.

For a collection, click it in the web library sidebar and copy the eight-character key after /collections/ in the URL. Leave the collection blank for the entire library.

### 2. Create a key and fill the form

Open <https://www.zotero.org/settings/keys> → Create new private key → name it LUNA references → allow read access to your Personal Library or required groups. Write access is unnecessary. Copy the generated key.

LUNA → Zotero → connection name → library type (group/personal) → numeric ID → optional collection key → API key in the token field → Save. The source URL is built from these IDs automatically. Add separate connections for other libraries/accounts.

### 3. Check the reference list

Collect now → project home → References → View list. Check authors, year, abstract and DOI/source links. In Materials combine the Reference category with project and tag filters.

After editing Zotero metadata, sync Zotero and collect again or wait for scheduled collection. Collection imports accessible bibliographic information; it does not grant access to paid paper PDFs.

## 23. Materials, classification and references

Find sources, inspect originals and correct classification and tags.

Open this guide online

### 1. Find a document

1. Open Materials and enter a title or body keyword in the basic search field.
2. Expand Advanced search and type to select grants, projects, researchers or tags. Add multiple selections and remove them with ×.
3. Choose the relevant kind: report, minutes, code, data, general material or reference, then search. Remove filters one by one if no results appear.
4. Open a title to inspect its body, source date, collection date and original link. Extracted text may differ from the original PDF or table layout.

### 2. Correct metadata

1. Open the classification and related-items area on a material page. Its connection creator, direct uploader or manager may save changes while able to read it; other readers may propose changes.
2. Correct its kind, assign grants, projects, researchers and tags, then save. Source paths drive automatic code classification, while manual classification is retained separately.
3. Review whether it should remain reference context. This makes it a background candidate for AI, not proof of work in the reporting period.
4. Return to search and verify the kind and tags. Assigning metadata does not grant other users access.

### 3. Use references

1. Open Project home → References for connected Zotero items.
2. Review titles, authors, dates, abstracts, DOIs and tags and follow DOI or original links.
3. Select Reference in Materials to search accessible bibliography across projects.
4. Edit metadata in Zotero, sync online and collect again. Zotero collection currently reads bibliography and abstracts, not automatically acquired PDF full text.

## Find materials in the grouped overview

The Materials page shows counts and recent items grouped by type. Expand an item to see its dates and a short excerpt, then choose Open material for its full copy and relationships. View all inside a group keeps your filters and that type; View all at the top shows all current matches. Advanced grant, project, researcher, tag, date and type filters remain available.

## 4. Generate tags and review them in place

1. Expand metadata editing on a material/report and choose Generate beside Tags (comma-separated).
2. The title, headings and body supply one to five keywords. Review the chips/field and status in place without returning to the top.
3. Use Keep existing tags to retain current tags while adding suggestions; when replacing them, review the proposed set.
4. Edit unsuitable tags and save metadata. Depending on your role, changes may become an approval proposal rather than an immediate edit; read the status message.
5. Tags neither rewrite the body nor grant access. Check other unsaved fields and the final saved/proposal state.

## Find materials with AI

1. Enter keywords in Materials and select Find with AI. The search opens on the right on desktop and fills the screen on mobile. Current date, grant, project, researcher, type and tag filters are retained.
2. Cards show excerpts, match reasons and dates; open the title to read the source. Metadata-only sources are labelled. Search uses no AI generation allowance.
3. Choose Ask about this source, or enter a summary/comparison question and choose Ask AI. This starts a private consultation using your personal AI allowance and the saved search boundary. Return to List and search filters to change the scope.
4. AI receives bounded excerpts, not every original, image or attachment. If there are no results, review the query and filters.

## Search terms and relevance

1. Enter several terms such as imagery eeg. Exact phrases rank first, then all terms, curated related terms and partial matches. Supported bilingual terms include EEG/뇌파 and imagery/심상. This does not translate arbitrary terms or find every semantically similar document.
2. Use "imagery eeg" for a required phrase, imagery OR eeg for either term, and eeg -animal to exclude a term. Quoted terms are not expanded. Use up to 200 characters and 16 terms. Suggested queries apply only when selected.

3. Choose Relevance or Default order. Date, grant, visibility, tag and current access restrictions stay in effect. Calendar items stay on their dates; the list below can rank by relevance. Announcements search only the published edition in the current language.

4. Basic search makes no external AI calls and consumes no personal or lab AI allowance. It searches the fields available on each page, such as a material title and stored body, without fetching inaccessible external files. Citation search uses titles only and preserves the conversation readers' access restrictions.

## Find and edit literature in Materials

1. Open References in Materials or search by title, author or DOI. References saved in a research space are discoverable in the same catalog; do not register another copy.

2. Select a card title to open its wide permanent detail page; View all opens the dedicated literature list. Translation, sharing, favorites and print are at the top, followed by authors, venue, publication date, abstract and source links. Original abstracts and AI summaries are separate. Generate an AI summary only with the explicit button or request.

3. Edit only the required bibliography fields. Materials and the space use the same reference ID and correction. Check the applicable research space; other spaces and Zotero originals are not silently overwritten.

4. Select references and choose BibTeX, RIS or CSL JSON under Export. Distinguish selected references, all search matches and the complete current scope. All includes every page in that scope, not just the visible rows.

5. Use File → Import in Zotero to read the exported file. This is a bibliographic snapshot, not automatic two-way synchronization. Unpinning does not delete the original material.

## Register references and request source storage in conversation

1. Tell native AI or ChatGPT the research space and the DOI, source URL or bibliography to register. Example: "Register these papers in this space, reuse existing references and do not generate new AI summaries."

2. Check the destination project/grant, audience, new/existing references and original attachments. The host must actually expose the required tools. Search-only tools cannot prove registration.

3. Check each returned material link and status. With a linked research work, the actual registration request and receipt also appear in its history. Ask for a destination if work context is unclear; never invent earlier conversation that was not delivered. Bibliography registration, pinning, source storage, text extraction and AI analysis are distinct. Reuse the same request ID when retrying to prevent duplicates.

4. Add a PDF later from the reference detail page. Active authorized participants can contribute files, preserving contributor and file versions. Identical files are not linked twice and earlier PDFs are not overwritten. Distinguish published papers, preprints and supplements. In conversation, identify the existing reference before attaching a PDF. Example: "Attach this PDF to the existing reference without creating

another paper.” Only delivered files or accessible originals can be stored.

5. A ZIP or download link generated in ChatGPT is not proof of platform storage. Check file delivery and the platform receipt. Registration, reading, pinning and exporting do not themselves generate AI summaries.

6. The tool sequence is `consultant_material_prepare` → `consultant_material_commit` → `consultant_material_status`. Review the destination and audience, then pass the unchanged `preview_token` and review. A cached host can use the existing `consultant_thread_prepare` save / materials envelope for the same service. If neither route is exposed, refresh the connection tool list.

7. A request accepts up to 50 items and 100 MB of combined attachments. Distinguish “register bibliography only” from “also store an accessible PDF.” A paper landing page is not reported as a saved PDF; supply an actual PDF file when needed.

## Understand storage, PDF and extraction states

1. Once registered, bibliography can be searched and edited before source collection finishes. Do not create duplicate references because earlier files are still processing.

2. Queued and running storage are distinct. Stored means a verified preservation receipt exists. An archived HTML page is not a PDF.

3. If the original is stored but text extraction fails, it remains downloadable. Combined source attachments may total 100 MB; safe text extraction supports up to 20 MB per file. Larger originals are preserved with an extraction-limit status. Scans, protected documents or unsupported formats may not yield readable text.

4. Retryable storage failures use bounded attempts and backoff. For persistent failures or access requirements, inspect the reason and supply an accessible file. Access and login restrictions are not bypassed.

5. A PDF downloadable from its reference is distinct from source archiving and text extraction. Linking a PDF does not prove external synchronization, extraction or AI analysis. AI analysis is a separate request with its own usage. Stored originals or extracted text do not mean AI analysis has completed.

## Bibliography and PDF visibility

New public-paper bibliography is visible by default to active signed-in platform users, not anonymous visitors. PDF originals, private analysis and research spaces retain their existing access. Reusing an existing private source does not publish it. Edit bibliographic fields and grant/project/researcher metadata on the permanent literature page. External Zotero originals are not overwritten.

## Find and attach PDFs from ChatGPT

Example: “Find references in this space that lack a PDF. Locate public originals, verify the title and DOI, and attach them to the existing references while preserving earlier files.” Tools import a public HTTPS PDF or an actually transferable ChatGPT file and verify bytes, size and hash. They never bypass login or paid access. Original preservation, text preparation and AI analysis are separate states; an unavailable file transfer is not a saved result.

## Compare literature and save to the same work

Example: “Compare these three papers for the Paper Goal and Literature Survey work. State the PDF version and actual coverage; tabulate model structure, data, metrics and limitations. Save the delivered analysis as another request/result pair in this same work and attach the comparison file.” The existing work ID and earlier results are preserved. An excerpt is never labelled a complete-paper reading. Specify explicitly when a new derived work is intended.

## File transfer and retry

Example: “Save the actual image and BibTeX files just created as outputs of this work. Report each saved status and download link.” Attach up to 50 files, combined 100 MB. Reuse the request key for retries; reselect files to renew expired transfer links. The file tool must actually be available in the ChatGPT host. If unavailable, transfer the files through Add sources on the platform. A filename or temporary link is not a preserved file.

## Edit, export and apply requests

Search, select required references or all references, then use the download icon beside search for BibTeX, RIS or CSL JSON. Example: “Correct the selected authors and DOI while preserving the original.” AI summaries run only after a button action or explicit request. Applying the refinement preview updates both the current composer and its local draft, ready to inspect before sending. Dashed graph edges indicate references; solid edges indicate research flow.

## Start paper analysis in the same work

In the existing work composer, use Add sources to select one reference or PDF. Enter a goal-specific question, choose Paper analysis and discovery in the Send to AI menu, and send. Source reading, backward/forward/topic search, registration, synthesis and manuscript drafts are checkpointed. Search uses the public paper title or your explicit public query, not private goals. Verified references and final recommendations are registered automatically; PDF collection is queued with actual storage status. Missing originals limit analysis to the abstract. Forward and similar searches cover the first 25 results each; request further searches separately. Paywalls and institutional login are not bypassed; unavailable originals retain a reason. Registration, original collection, AI analysis and author review are separate states.

## Goals, notifications and collaborative review

Use work menu → Research conditions to edit the question, venue, deadline and constraints. Expand Collaborative guidance to opt into a reviewed goal snapshot and next-action notifications. Only the published goal version is shared; private edits do not silently replace it. Notifications open the same suggestion ID used by the composer and do not execute AI. Deferred or dismissed suggestions are not resurrected. Before joint use, check work sharing, access to originals and the shared goal separately. Approved methods also require use and provider-processing permission. Participants use the research team destination for follow-up analysis. Private results are not published automatically. A researcher confirms submission or goal completion.

## Improve v1 into v2

Use work menu → Know-how settings to select exact approved versions. The default direction and relevant common/activity guidance are pinned; four stages are not required for every work. Existing work-only know-how remains available in collapsed history. Select currently accessible, non-deleted request/result evidence and distinguish paper-specific fixes from reusable improvements. Compare against the actual active v1 using matched input/model/work conditions, held-out input and human review receipts before considering v2 promotion. Promotion changes future defaults and is reversible; it does not rewrite existing results. The built-in method before initialization is distinct from a reviewed and evaluated v1. Method versions, manuscript revisions and model IDs are distinct. Output-only use, source reading and derivation are separate permissions.

## Delete and restore request/result pairs

The author or an authorized space manager can move an accessible request/result pair to trash through its menu. Removed content is excluded from new analysis, method improvement and downloads; other researchers' follow-ups and existing quotations remain. Restore through Deleted requests and results; this never restarts AI. Management does not grant access to another person's private conversation. External copies and provider costs already incurred are unaffected.

## ChatGPT, consultation, Overleaf and costs

Continue in ChatGPT preserves the same work ID and selected evidence, saving only delivered results. Survey request/status/resume/external, method reads/improvements and pair delete/restore reuse consultant\_thread\_prepare/commit reviews. Independent host analysis never claims to have executed a protected method. Core instructions are excluded from ordinary results, history and exports. Provider processing permission is reviewed separately; this does not guarantee that outputs cannot reveal or imply knowledge. Download the draft and BibTeX, review them, then incorporate them in Overleaf; there is no automatic bidirectional sync. Per-request, daily and provider limits pause work while preserving completed stages. Inspect saved stages and usage before explicitly resuming. Team consultation tools currently support requests and status; use the web for resuming and method management. Host tool availability, new server PDF collection, AI quality and Lightsail latency require separate verification.

## Keep only the know-how you choose from a discussion

In the same work composer, ask: “Separate the literature-analysis and manuscript-writing know-how from this discussion. Propose only the changes to each existing method.” Only delivered, saved requests and results are evidence. External ChatGPT or Codex can prepare structured proposals using returned evidence IDs and current permissions, then use the existing reviewed save contract. On the proposal page, select only the drafts you want to keep and review the save. Unselected proposals and active versions remain unchanged. Preparing or saving a draft does not validate quality or approve v2.

## Apply several methods together

Use the applied know-how count below the work title, or work menu → Know-how settings, to manage defaults. No empty badge appears for zero items. Search names, handles or activities and review exact versions and roles before saving. Add to a work on the independent detail page lets you choose a space and work and pin an exact version. Type @ in a follow-up, new research request or personal AI consultation to search researchers and know-how together. Choosing know-how adds it only to that request without changing defaults. [[ source citations keep their existing behavior. Multiple selections never auto-merge assets or adopt newer versions. Web, consultation and ChatGPT handoff check the same IDs and current permissions.

## Split, merge, originals and shared evidence

One discussion can produce independent proposals that reference the same segment as shared evidence. Drafts start privately. Separate paper-specific corrections and unverified scientific conclusions from reusable guidance. Change organization prepares a split or merge request in the existing composer for review and sending. A merge proposes a new asset linked to exact parent versions. It does not delete or overwrite the versions, history or active state of assets A and B. Changing how a proposal is organized requires a newly reviewed proposal, not an unreviewed browser-side rewrite.

## Version inheritance and evaluation

Derived work inherits the exact selected versions of independent assets. For existing work-specific methods, the collapsed Work know-how controls retain version and candidate selection in the original work. Review generalized instructions and create a separate asset before reuse in other work. Execution applies modules relevant to the current stage, such as literature, synthesis, writing or review, while retaining links to other stages. New versions are not adopted automatically. Initial approval reviews the actual instructions and their generalization; it does not claim measured quality improvement. Improvement candidates require actual runs with matched questions, inputs, models and effort, held-out cases and human review before adoption. A candidate trial is a one-off selection that does not change the default and uses the normal AI allowance.

## Protection, storage and execution boundaries

An independent detail page shows introduction, ID, handle, authors, usage, current permissions and versions. Editing the introduction or image link does not change execution guidance; image access is checked separately. The default direction is protected guidance. Write a draft or propose a direction improvement, review it and select which drafts to keep. Saving an introduction or candidate does not change the active version. Ordinary catalogs and proposal cards exclude protected instructions. Source review requires current permission and AI-provider processing permission is separate. Reads, selections and structured proposal saves do not run AI, while sending analysis or reorganization requests may use your allowance. Independent external analysis does not prove protected-service execution. Instructions belong in encrypted backups, not code or static files. AI quality, actual ChatGPT tool availability and server speed require separate validation.

## 24. Upload materials and choose cloud files

Upload files, import selected cloud files and request optional analysis.

Open this guide online

### Register a file as research material

1. Open Materials, then Upload material.
2. Choose a file, up to 100 MB. Research-space attachments allow up to 50 files, 100 MB combined. Use an external connection for larger files. The browser sends chunks and displays progress.
3. Optionally enter a title; otherwise the filename is used. Choose a lab and at least one grant or project. Add researchers and a memo; @mentions and [[citations]] work in the memo.
4. Select the checkbox under 3. Optional AI analysis if needed (off by default), then select Register material. Keep the tab open during transfer. Reaching the material page means transfer was accepted; Drive storage and text extraction run in the background; AI runs only when requested.
5. Download original becomes available after verified storage. For failures, read the status and retry. Storage retries reuse a file ID to avoid duplicate originals.
6. Administrator prerequisite: enable the existing Google Drive archive with a nonpublic writable folder and archive account. A read-only source connection is insufficient for storage.

### Audience and direct-upload search

1. Current readers and managers of a selected grant or project can view the material. Merely tagging a researcher does not grant access.
2. The author or a current manager can edit metadata. Changing grants/projects changes subsequent viewing, search and download access. At least one scope is required.
3. In Materials search, choose Directly registered materials under Registration method to filter files, links and cloud imports.
4. Platform access changes do not recall downloaded copies or change Google Drive sharing permissions. Administrators must also manage the archive folder permissions.

### Notion exports, public links and YouTube

1. Export private Notion pages as PDF, Markdown or CSV, then upload the file. ZIP exports containing multiple pages and attachments are supported.
2. Extracted headings, text and tables are displayed as a material page. Supported documents inside a ZIP are read, included images can be expanded, and the original ZIP remains downloadable. Extraction size

and format limits are separate from storage limits; a stored original remains downloadable if extraction fails.

3. For public web or Notion pages, paste the public URL instead of selecting a file. Only accessible HTML that permits crawling is collected. For login- or JavaScript-dependent pages, upload an export.

4. YouTube URLs provide an embedded player. This does not collect the original video or captions or claim full-video AI analysis.

## Choose from Google Drive or OneDrive

1. Open Materials → Upload material → 1. Choose material → Choose from cloud. Select Connect Google Drive or Connect OneDrive, choose your own account and allow reading. No project connection or predetermined folder is required.

2. Browse folders in My Drive or search All files and shared materials in Google Drive. For OneDrive, browse and search your drive. Lists show 50 entries at a time, with Next files and parent-folder breadcrumbs.

3. Select a file. Google Docs, Sheets and Slides export to DOCX, XLSX and PPTX. Files above 100 MB, unavailable downloads and unsupported formats display their reason and cannot be selected. Provider access and organization policies still apply.

4. Back in Upload material, check the filename, lab and audience before registering. Starting from a grant/project connection page preselects that scope. Only the chosen file is copied to the lab Google Drive archive; the source is not moved or deleted.

5. Change account/reconnect lets you choose another account. Disconnect stops the app using this personal credential; existing collection connections and imported materials remain. File selections expire after 30 minutes or when the account changes/disconnects.

6. A platform administrator performs the initial Microsoft setup once. Administrator setup and step-by-step guide covers Entra registration, personal+organizational accounts, Web callback, client ID, secret Value, and delegated Files.Read/offline\_access. Researchers only sign in. Existing OneDrive source-link connections remain available; selected-file imports are separate from scheduled collection.

7. Google reuses the lab OAuth app configuration. Testing mode requires test-user registration; Drive reading consent, Google verification and organizational policies still apply. Each user completes their own consent.

## Analysis and preview coverage

1. Text is extracted from TXT, Markdown, CSV, DOCX, HWP5, HWPX, PDF, PPTX and XLSX, among others. PDF pages, images and embedded Office raster images are analyzed in units. Protected/encrypted HWP files are not bypassed.

2. Change the PDF preview page number, expand image previews or play MP4/WebM/MOV in the video player. Use the separate button for original downloads.
3. Video AI examines frames sampled every 30 seconds within the first 10 minutes. Audio transcription and exhaustive frame analysis are not included.
4. ZIP limits are 100 entries, 100 MB total expanded, 20 MB per entry. Supported inner documents are analyzed; nested ZIPs, encrypted and executable files are omitted. Overall limits include 300 analysis units and 200,000 displayed text characters. Inspect File list and analysis coverage for omissions.
5. The uploader or a platform administrator can select Start/Resume AI analysis on the material page. Active lab settings and the requester's current AI permission are required. Each page, image or text segment uses one attempt from the uploader's daily allowance, even for administrator requests; failed attempts also count. Quota exhaustion resumes the next day. Storage and text extraction alone do not use AI allowance. Results can inform AI consultation, but remain interpretations separate from the source; verify important numbers and conclusions against originals.

## Check and recover original-file errors

Use Check original on the material detail to see the current readability check and time. Historical storage success is not a successful download now. Changing the default archive folder does not change the saved file and folder receipt for older originals. Reconnect the existing archive account for expired authorization and follow the displayed code for moved, deleted or inaccessible files. Do not make files public or substitute similarly named files. Text extraction, attachment preservation and image or handwriting interpretation are separate results.

## Find and edit literature in Materials

1. Open References in Materials or search by title, author or DOI. References saved in a research space are discoverable in the same catalog; do not register another copy.
2. Select a card title to open its wide permanent detail page; View all opens the dedicated literature list. Translation, sharing, favorites and print are at the top, followed by authors, venue, publication date, abstract and source links. Original abstracts and AI summaries are separate. Generate an AI summary only with the explicit button or request.
3. Edit only the required bibliography fields. Materials and the space use the same reference ID and correction. Check the applicable research space; other spaces and Zotero originals are not silently overwritten.
4. Select references and choose BibTeX, RIS or CSL JSON under Export. Distinguish selected references, all search matches and the complete current scope. All includes every page in that scope, not just the visible rows.
5. Use File → Import in Zotero to read the exported file. This is a bibliographic snapshot, not automatic two-way synchronization. Unpinning does not delete the original material.

## Register references and request source storage in conversation

1. Tell native AI or ChatGPT the research space and the DOI, source URL or bibliography to register. Example: “Register these papers in this space, reuse existing references and do not generate new AI summaries.”
2. Check the destination project/grant, audience, new/existing references and original attachments. The host must actually expose the required tools. Search-only tools cannot prove registration.
3. Check each returned material link and status. With a linked research work, the actual registration request and receipt also appear in its history. Ask for a destination if work context is unclear; never invent earlier conversation that was not delivered. Bibliography registration, pinning, source storage, text extraction and AI analysis are distinct. Reuse the same request ID when retrying to prevent duplicates.
4. Add a PDF later from the reference detail page. Active authorized participants can contribute files, preserving contributor and file versions. Identical files are not linked twice and earlier PDFs are not overwritten. Distinguish published papers, preprints and supplements. In conversation, identify the existing reference before attaching a PDF. Example: “Attach this PDF to the existing reference without creating another paper.” Only delivered files or accessible originals can be stored.
5. A ZIP or download link generated in ChatGPT is not proof of platform storage. Check file delivery and the platform receipt. Registration, reading, pinning and exporting do not themselves generate AI summaries.
6. The tool sequence is `consultant_material_prepare` → `consultant_material_commit` → `consultant_material_status`. Review the destination and audience, then pass the unchanged `preview_token` and review. A cached host can use the existing `consultant_thread_prepare` save / materials envelope for the same service. If neither route is exposed, refresh the connection tool list.
7. A request accepts up to 50 items and 100 MB of combined attachments. Distinguish “register bibliography only” from “also store an accessible PDF.” A paper landing page is not reported as a saved PDF; supply an actual PDF file when needed.

## Understand storage, PDF and extraction states

1. Once registered, bibliography can be searched and edited before source collection finishes. Do not create duplicate references because earlier files are still processing.
2. Queued and running storage are distinct. Stored means a verified preservation receipt exists. An archived HTML page is not a PDF.
3. If the original is stored but text extraction fails, it remains downloadable. Combined source attachments may total 100 MB; safe text extraction supports up to 20 MB per file. Larger originals are preserved with an extraction-limit status. Scans, protected documents or unsupported formats may not yield readable text.

4. Retryable storage failures use bounded attempts and backoff. For persistent failures or access requirements, inspect the reason and supply an accessible file. Access and login restrictions are not bypassed.

5. A PDF downloadable from its reference is distinct from source archiving and text extraction. Linking a PDF does not prove external synchronization, extraction or AI analysis. AI analysis is a separate request with its own usage. Stored originals or extracted text do not mean AI analysis has completed.

## Bibliography and PDF visibility

New public-paper bibliography is visible by default to active signed-in platform users, not anonymous visitors. PDF originals, private analysis and research spaces retain their existing access. Reusing an existing private source does not publish it. Edit bibliographic fields and grant/project/researcher metadata on the permanent literature page. External Zotero originals are not overwritten.

## Find and attach PDFs from ChatGPT

Example: “Find references in this space that lack a PDF. Locate public originals, verify the title and DOI, and attach them to the existing references while preserving earlier files.” Tools import a public HTTPS PDF or an actually transferable ChatGPT file and verify bytes, size and hash. They never bypass login or paid access. Original preservation, text preparation and AI analysis are separate states; an unavailable file transfer is not a saved result.

## Compare literature and save to the same work

Example: “Compare these three papers for the Paper Goal and Literature Survey work. State the PDF version and actual coverage; tabulate model structure, data, metrics and limitations. Save the delivered analysis as another request/result pair in this same work and attach the comparison file.” The existing work ID and earlier results are preserved. An excerpt is never labelled a complete-paper reading. Specify explicitly when a new derived work is intended.

## File transfer and retry

Example: “Save the actual image and BibTeX files just created as outputs of this work. Report each saved status and download link.” Attach up to 50 files, combined 100 MB. Reuse the request key for retries; reselect files to renew expired transfer links. The file tool must actually be available in the ChatGPT host. If unavailable, transfer the files through Add sources on the platform. A filename or temporary link is not a preserved file.

## Edit, export and apply requests

Search, select required references or all references, then use the download icon beside search for BibTeX, RIS or CSL JSON. Example: “Correct the selected authors and DOI while preserving the original.” AI summaries run only after a button action or explicit request. Applying the refinement preview updates both the current composer and its local draft, ready to inspect before sending. Dashed graph edges indicate references; solid edges indicate research flow.

## Start paper analysis in the same work

In the existing work composer, use Add sources to select one reference or PDF. Enter a goal-specific question, choose Paper analysis and discovery in the Send to AI menu, and send. Source reading, backward/forward/topic search, registration, synthesis and manuscript drafts are checkpointed. Search uses the public paper title or your explicit public query, not private goals. Verified references and final recommendations are registered automatically; PDF collection is queued with actual storage status. Missing originals limit analysis to the abstract. Forward and similar searches cover the first 25 results each; request further searches separately. Paywalls and institutional login are not bypassed; unavailable originals retain a reason. Registration, original collection, AI analysis and author review are separate states.

## Goals, notifications and collaborative review

Use work menu → Research conditions to edit the question, venue, deadline and constraints. Expand Collaborative guidance to opt into a reviewed goal snapshot and next-action notifications. Only the published goal version is shared; private edits do not silently replace it. Notifications open the same suggestion ID used by the composer and do not execute AI. Deferred or dismissed suggestions are not resurrected. Before joint use, check work sharing, access to originals and the shared goal separately. Approved methods also require use and provider-processing permission. Participants use the research team destination for follow-up analysis. Private results are not published automatically. A researcher confirms submission or goal completion.

## Improve v1 into v2

Use work menu → Know-how settings to select exact approved versions. The default direction and relevant common/activity guidance are pinned; four stages are not required for every work. Existing work-only know-how remains available in collapsed history. Select currently accessible, non-deleted request/result evidence and distinguish paper-specific fixes from reusable improvements. Compare against the actual active v1 using matched input/model/work conditions, held-out input and human review receipts before considering v2 promotion. Promotion changes future defaults and is reversible; it does not rewrite existing results. The built-in method before initialization is distinct from a reviewed and evaluated v1. Method versions, manuscript revisions and model IDs are distinct. Output-only use, source reading and derivation are separate permissions.

## Delete and restore request/result pairs

The author or an authorized space manager can move an accessible request/result pair to trash through its menu. Removed content is excluded from new analysis, method improvement and downloads; other researchers' follow-ups and existing quotations remain. Restore through Deleted requests and results; this never restarts AI. Management does not grant access to another person's private conversation. External copies and provider costs already incurred are unaffected.

## ChatGPT, consultation, Overleaf and costs

Continue in ChatGPT preserves the same work ID and selected evidence, saving only delivered results. Survey request/status/resume/external, method reads/improvements and pair delete/restore reuse consultant\_thread\_prepare/commit reviews. Independent host analysis never claims to have executed a protected method. Core instructions are excluded from ordinary results, history and exports. Provider processing permission is reviewed separately; this does not guarantee that outputs cannot reveal or imply knowledge. Download the draft and BibTeX, review them, then incorporate them in Overleaf; there is no automatic bidirectional sync. Per-request, daily and provider limits pause work while preserving completed stages. Inspect saved stages and usage before explicitly resuming. Team consultation tools currently support requests and status; use the web for resuming and method management. Host tool availability, new server PDF collection, AI quality and Lightsail latency require separate verification.

## Keep only the know-how you choose from a discussion

In the same work composer, ask: “Separate the literature-analysis and manuscript-writing know-how from this discussion. Propose only the changes to each existing method.” Only delivered, saved requests and results are evidence. External ChatGPT or Codex can prepare structured proposals using returned evidence IDs and current permissions, then use the existing reviewed save contract. On the proposal page, select only the drafts you want to keep and review the save. Unselected proposals and active versions remain unchanged. Preparing or saving a draft does not validate quality or approve v2.

## Apply several methods together

Use the applied know-how count below the work title, or work menu → Know-how settings, to manage defaults. No empty badge appears for zero items. Search names, handles or activities and review exact versions and roles before saving. Add to a work on the independent detail page lets you choose a space and work and pin an exact version. Type @ in a follow-up, new research request or personal AI consultation to search researchers and know-how together. Choosing know-how adds it only to that request without changing defaults. [[ source citations keep their existing behavior. Multiple selections never auto-merge assets or adopt newer versions. Web, consultation and ChatGPT handoff check the same IDs and current permissions.

## Split, merge, originals and shared evidence

One discussion can produce independent proposals that reference the same segment as shared evidence. Drafts start privately. Separate paper-specific corrections and unverified scientific conclusions from reusable guidance. Change organization prepares a split or merge request in the existing composer for review and sending. A merge proposes a new asset linked to exact parent versions. It does not delete or overwrite the versions, history or active state of assets A and B. Changing how a proposal is organized requires a newly reviewed proposal, not an unreviewed browser-side rewrite.

## Version inheritance and evaluation

Derived work inherits the exact selected versions of independent assets. For existing work-specific methods, the collapsed Work know-how controls retain version and candidate selection in the original work. Review generalized instructions and create a separate asset before reuse in other work. Execution applies modules relevant to the current stage, such as literature, synthesis, writing or review, while retaining links to other stages. New versions are not adopted automatically. Initial approval reviews the actual instructions and their generalization; it does not claim measured quality improvement. Improvement candidates require actual runs with matched questions, inputs, models and effort, held-out cases and human review before adoption. A candidate trial is a one-off selection that does not change the default and uses the normal AI allowance.

## Protection, storage and execution boundaries

An independent detail page shows introduction, ID, handle, authors, usage, current permissions and versions. Editing the introduction or image link does not change execution guidance; image access is checked separately. The default direction is protected guidance. Write a draft or propose a direction improvement, review it and select which drafts to keep. Saving an introduction or candidate does not change the active version. Ordinary catalogs and proposal cards exclude protected instructions. Source review requires current permission and AI-provider processing permission is separate. Reads, selections and structured proposal saves do not run AI, while sending analysis or reorganization requests may use your allowance. Independent external analysis does not prove protected-service execution. Instructions belong in encrypted backups, not code or static files. AI quality, actual ChatGPT tool availability and server speed require separate validation.

## 25. Change requests and trash

Submit and review changes, handle conflicts and restore items.

[Open this guide online](#)

### Submitting a proposal

1. Open a material or AI activity report you can currently read.
2. Expand Edit metadata within the metadata section (Edit classification and related items for materials).
3. Search and select grants, projects and researchers, then enter tags. For materials you can propose a category or AI reference flag; for reports you can also propose a title and memo.
4. Optionally explain why, then press Propose changes. Managers see Save in the same editor. Existing values remain until approval. Source content and the original report generation conditions are unchanged.
5. Check the proposal page for its status or withdraw your pending proposal. My proposals on the same material lets you return to previous submissions.

### Reviewing as a manager

1. Open Profile → Settings → Change requests. Track sent and reviewable requests together. Connection creators, direct uploaders and existing managers may approve or reject currently readable materials. Researcher tags do not confer approval rights.
2. Filter by status and search the material title, proposer or reason. Only proposals you may manage are shown.
3. Check proposals and choose Approve selected, or Approve all search results. Each proposal is checked for current permissions/conflicts and skipped counts are shown. Open a proposal to reject it individually.

Compare Before and Proposed and enter a review note.

4. Approve and apply, or reject. The reviewer, time, note and before/after values are retained.
5. Approval stops if metadata or participation permissions changed. Review the current material and request a new proposal. Adding a report audience also requires management rights over the added grants/projects.

Submission and decision notifications follow existing system-notification preferences. Managers can still edit metadata directly.

## When a report is visible but its evidence is not

Report sharing and source access are separate. A grant member who has not joined a child project may read its shared report without access to source materials collected by that project. Request project membership or contact its manager. Disabled or changed source connections can also hide evidence. Editing metadata or assigning a researcher does not grant source access.

## Deletion requests and your trash

1. Open a material, report, grant, project or conversation and select Move to trash or Request deletion. You can directly trash items you created or currently manage. Assigning yourself in researcher metadata does not make you the creator.
2. Without deletion permission, optionally explain the reason, check the confirmation and submit. The item remains until approval.
3. Open Profile → Settings → Change requests to track status and review notes or withdraw a pending request. Submission and decision alerts follow platform notification preferences.
4. Managers open Profile → Settings → Change requests and select Deletion request as the type. Review the target and reason, then approve or reject. Bulk approval includes deletion requests: verify the type and targets. Changed permissions or item state block approval.
5. Open Profile → Settings → Trash. Search or filter by type. Restore individually, selected items or all search results. Restore parent grants, projects and channels first. Restored external connections require authentication and collection settings to be checked separately.
6. For permanent deletion, review all affected items and enter your password. Deletion is blocked if it would affect another person's private AI conversation or data outside your authority. Contact a platform administrator in that case. Deleted collected items stay suppressed during collection for the same connection version; reconfiguring a connection may create a new collection scope.
7. Restoring a conversation preserves its history and earlier individual deletions. This feature does not erase Google, GitHub or Notion originals or external backups. Only the authenticated calendar owner may explicitly delete a Google original from the separate event deletion page. Platform administrators delete the platform copy only.

## Actual dates and AI review

1. Save Actual material date and Date evidence in material metadata. Provider timestamps remain separate; clearing the date restores the existing basis.
2. AI metadata review proposes dates, categories and tags from available collected evidence. It does not invent facts from unread content or apply proposals before approval.
3. See Material AI analysis for execution, pause/resume and budgets. Both metadata and content analysis for a fixed initial cohort use a separate allowance; this exemption does not extend to every new material.

4. Completed analysis, pending proposals and approved changes are different states. Compare old and proposed values in Change requests before approving or rejecting.

### Automatically populated fields and verification

1. Inspect the collected document's grant/project scope and attributed authors, assignees and researchers.
2. Calendar attendees may link to an unambiguous connected account in the lab. An attendee listing does not prove actual attendance.
3. Correct missing or incorrect links through editing or a proposal. Automatic researcher association does not prove sole authorship or responsibility.
4. Review manual and automatic values before saving. Authorized approvers review proposals and can open the document from the request alert.

### Choose review request notifications

1. Open Change request notification settings from the review inbox. Expand the initially collapsed Announcements, requests and operations group in notification preferences.
2. Platform delivery is selected and Digest is clear by default for Metadata/change requests and review results. Metadata and deletion requests you can review appear in your inbox without waiting for a digest.
3. To opt out, clear Platform and save. Email requires both the overall email switch and this category. Device push also requires its overall switch and browser permission.
4. To batch incoming requests, select Digest on that row and choose an interval, then save. Incoming requests and review results follow the same digest selection. Administrator-disabled categories remain blocked.

## 26. Manage material AI analysis

Manage analysis progress, pause/resume and daily allowances.

[Open this guide online](#)

### Choose analysis and inspect progress

1. Platform administrators open Lab management → Material AI analysis. Select metadata, content or both and expand categories to check the intended scope.
2. Metadata analysis creates proposals for dates, categories and tags; approval is required. Content analysis creates notes about text, pages and images.
3. Distinguish completed, running, queued and failed items. Inspect status, errors and source dates. Preserved historical results excluded from current materials after collection failures are marked separately.
4. Pause/resume affects that analysis track. An in-flight response may finish; resuming does not regenerate every completed result.

### Normal allowances and the fixed initial cohort

1. Check the current shared analysis and personal AI allowances. Normal content analysis uses attempts per page, image or text segment, including failed attempts where applicable.
2. A fixed initial cohort has a separate daily budget for both metadata and content. The default and maximum is 100 materials; zero pauses it. Multiple pages of one material count as one material that day.
3. New materials and newly observed versions do not inherit this exemption. Material counts are not exact token/cost limits or guaranteed completion counts.
4. Quota-bound work resumes through the worker after the Seoul date changes. Permissions, source changes, configuration and pause controls still apply.

### Separate registration, consultation and backup

1. AI analysis of direct uploads is optional and off by default. The uploader or a platform administrator requests it; normal items use the uploader's allowance.
2. Notes can be reused as consultation evidence candidates, but every file is not read in full for each question. Inspect omissions and partial coverage.
3. Check preservation and Drive-copy status separately from completed analysis. Shared research copies and private consultation storage have different boundaries.
4. Analysis is AI interpretation. Compare source dates, figures and numbers and approve only appropriate metadata proposals.

## Personal AI and shared analysis allowance policy

1. Settings → AI usage and remaining allowances is available to every user for personal usage, remaining attempts and their labs' shared analysis budgets. Platform administrators manage personal allowances.
2. AI help, personal AI consultations, announcement/document translation, reports and magazines use personal allowances. Selecting a lab in AI help chooses its execution configuration and does not spend its analysis budget. The legacy separate 10-attempt cap no longer applies.
3. Material AI analysis → Daily analysis attempt limit for regular materials is the sole source of the shared analysis budget. Saving 500 means 500 attempts per day across metadata, file, text and automatic conversation analysis. Existing personal allowances also apply to regular analysis; direct uploads are charged to the uploader. Regular analysis therefore also waits if the personal allowance runs out first.
4. A fixed initial cohort uses a separate limit of up to 100 materials per day, outside personal and regular analysis allowances. New materials and versions do not inherit that exemption.
5. Daily usage resets at midnight Seoul time. Failed execution attempts count, but permission or allowance rejection does not. Queued and duplicate requests are counted only when an execution starts. Long translations and multipage analysis may need multiple executions. On-device browser translation and handbook search use no AI attempts.
6. Earlier failures remain historical records. Ask again if personal attempts are currently available. Provider subscription/rate limits are separate from app allowances; AI permission and an enabled lab configuration remain required.

## AI management: settings in one place

Personal settings → Personal AI settings always shows only your effective policy, translation and usage, even for administrators. Lab management → Platform AI management is a separate administrator page for execution/resume policy for all users, platform statistics, user permissions and prices. Edit notifications on the dedicated notification settings page. A request budget is separate from the daily allowance. New work snapshots saved defaults; only platform administrators can change the central policy. Administrator defaults are 50 attempts and 3000 seconds per call, with no overall elapsed-time deadline. Deliberate resume applies current policy without resetting usage. Reductions cap subsequent calls of existing requests. Completed stages are reused.

## Attempts, tokens and estimated costs

Attempts are not invoices. New executions record feature, user, grant, project and requested model, with daily/weekly/monthly/yearly or custom-date views. Filtered totals above and below the table cover calls, input/output tokens and estimated USD across the complete current filter, not just displayed rows. Check measured-call coverage against all calls for each metric. Unknown tokens and costs remain unmeasured; historical gaps are not zero. Administrator-entered rates with a source/date apply to estimates for subsequent executions. Subscription fees and invoices are separate.

## Paused work, source access and attachments

Analysis detail preserves the request and work status. If source access is unavailable, affected output is hidden with connection/permission guidance; it is not deleted. Recheck after access is restored. Request history distinguishes loading, empty, failed and retry states. PDF page failures preserve successful page extraction and show causes such as encryption, timeout or pages without text. Extracted text, visual inspection and experimental validation are distinct.

## Personal settings through ChatGPT and AI consultation

Ask “Show my remaining AI allowance and reset time”, “Show my contribution score”, or “Show my administrator-managed execution policy”. Shared account tools read current settings and prepare an exact before/change review. Execution policy changes require current platform administrator authority and explicitly review the effect on all users. Personal overrides are rejected. Other settings remain personal; sharing rights and publication remain separate. Tools: `consultant_account_read`, `prepare`, `commit`, `usage`, `contributions` and `diagnose`. Help AI can read current personal settings and usage. Continue to work/settings consultation carries the question into private full consultation; subsequent reasoning uses personal AI allowance and changes require review. If the ChatGPT host does not expose a tool, check the connection and granted scopes.

## Repository consultation, isolated checks and integration

In GitHub code consultation choose architecture/paper mapping, README, function/class documentation or commit/research changes. The inspected file coverage and up to 20 commit messages are explicit; execution is never inferred. ZIPs include inspected source plus changes under `standalone/`, `INTEGRATION.md`, patches and check receipts. Run toy examples and tests against the real generated API in an isolated environment before integrating on a separate branch at the pinned baseline. Zero tests or missing dependencies do not establish verification.

## Detailed filters and research notifications

Detailed filters search authorized users, grants and projects even without recorded usage. Select all translation or individual features including help. Unassigned means missing attribution; unmeasured tokens are not zero. The top area shows active work, issues and unseen results. Mark as seen acknowledges a notification, not scientific approval. Cancelled and reviewed work remains in the goal/code history and archive. Saved materials are preserved.

## Analysis continuity and ChatGPT

New strategy analyses default to prior reports, feedback, saved follow-up decisions and recent discussions for the same goal. Originals remain available; bounded excerpts and omissions are recorded for model input. Independent analysis excludes prior analysis context. Revoked evidence is never reused. Code consultation from a report, section or action retains source provenance. ChatGPT resumes the same request ID and checkpoints; completed work uses follow-up analysis. analysis with `status=true` returns small metadata. Review `strategy_prepare` acknowledge with `kind/id/version` before commit. If the host exposes an old limit or missing operations, refresh the connected tools and read the central policy.

## Synchronizing external discussion and code results

Each code workspace has one representative notice across preparation, AI and validation steps. Dismiss preserves failures and does not mark them successful. Link a strategy report under Sources and links; the report also links back to derived code. External results preserve the question, answer, files, provenance, versions and downloads. Ask ChatGPT or Codex to continue this LUNA workspace and save the answer and files to the same workspace. Only a ready storage receipt confirms synchronization. Conversations or temporary files never delivered to LUNA are not automatically imported. You can also paste an answer and select files in the platform. Conversation text is limited to 120 KB, files to 100 files and 8 MB total; retry with identical files and request identity to resume. Files are not executed; model and test claims are producer-reported. Revisions preserve previous results. Storage uses no new AI call but consumes storage and backup capacity. Encrypted rejected responses may be revalidated without AI for up to 24 hours; invalid responses remain failures. Execution/resume settings start collapsed.

## Calendar ranges and official reference prices

Today/this week/this month/this year set date bounds in Asia/Seoul. Choose chart granularity separately in advanced filters. Administrators receive OpenAI Standard reference defaults checked on 2026-10-05. Custom rates and historical estimates are preserved. Long context, cache writes and processing tiers are distinguished; unknown tiers use an API-equivalent estimate, not subscription billing.

## Active work, archives and deletion

The live strip shows queued/running requests only. Analysis archives separate attention, completed, cancelled/failed and trash; saved outputs remain downloadable. Cancel before moving active work to trash; restore never resumes automatically. Permanent deletion requires an exact scope review and current permission, and rejects references from later consultations/code. Original attachments and usage totals remain. Old encrypted backups are not immediately erased; reapply the latest deletion ledger after restoration. ChatGPT uses `consultant_strategy_prepare analysis_trash/analysis_restore/analysis_purge` and `commit`.

## Complete literature and report history

Goal settings precede research context. Paginate all literature/search observations and reports. Deduplicated literature and per-search snapshots are separate; legacy observations have unknown original search batches. Only the final analysis stage creates a final report; partial results remain on the request. Compare two saved reports without AI. ChatGPT `strategy_read collection=reports/literature/observations` with `goal,page` reads the same lists.

## Start repository consultation

Choose a project, repository and purpose using Start GitHub repository consultation on the code landing page. Continue an existing analysis where available. Consultation is expanded on the detail page and requires a completed repository read. Connection/permission problems link to connection settings. The current bound of 24 text files/85KB does not establish whole-repository verification.

## Evidence access and analysis recovery

Analysis detail shows unavailable evidence counts and Google reauthorization status. Open the linked connection, personally authorize Google, wait for collection, then check access again. Authorization alone never restores historical access. Reuse saved stages only after a new authorized collection confirms the same file, version and content. Changed content or revoked access keeps old results preserved but restricted. Review a new analysis with accessible evidence to create a separate four-stage request after reviewing the original question, budget and timeout. Blocked source text and old outputs are not copied; new calls consume the existing personal allowance. Preview/status reads never run AI. ChatGPT uses `strategy_read` analysis diagnosis and the exact `strategy_prepare` analysis\_recover / `strategy_commit` review contract.

## Research daily limits and the overall allowance

Use Manage limits in Today's research to open Daily usage limits in Platform AI management.

Administrators choose Follow overall AI allowance or Set a separate cap. Administrators inherit by default; the existing regular-user default remains 5. Inheritance is not extra credit: research shares the personal pool with translation, help and other AI. A separate cap of 0 disables research AI.

Manage my overall daily AI allowance opens per-user overall limits. Attempted runs, queued reservations, completed and failed runs and reset times are distinct.

The overall pool resets at Seoul midnight; a separate research cap resets at midnight in the research timezone. Per-request budgets and per-call timeouts remain separate execution settings. Raising limits never clears history or usage, or automatically retries failures.

In ChatGPT or AI consultation, read `research_usage` via `consultant_account_read`. Only administrators can review and commit execution changes to `research_daily_limit`/`admin_research_daily_limit`; null inherits the overall pool.

## 27. Calendar and Google event creation

Create events, invite researchers and find recordings and notes.

Open this guide online

### 1. Prepare a calendar

A platform/lab administrator opens Profile → Lab management → Shared calendar settings. In Google Calendar, My calendars → ⋮ next to the calendar → Settings and sharing → Integrate calendar → copy Calendar ID into LUNA and save. Use the actual email-shaped Calendar ID; do not enter the alias primary. Connect a Google account that can write to this calendar. Other researchers should ask the calendar administrator.

### 2. Create an event and Meet

Calendar → Create event → Connect Google event creation permission. Choose an account and allow calendar event access. Read-only authorization cannot create events. Return to the form and choose the calendar, title, start and end. New dates default to today, so check and adjust them. Times are in Seoul time; all-day events include the chosen end date. Enable Create Google Meet to request a meeting link. After saving, use Copy link in the Google Meet card to copy just the meeting address. LUNA checks briefly while it is being created; Check meeting link only reads the existing event. In Invite researchers, search by name and select an email chip. The invitation address comes from Edit my profile → Google Calendar invitation email, or the profile email if left blank. Enter external guests under Additional invitation emails. Duplicate addresses are invited once. In Memo use [[title to find material or @username to select researchers. Select multiple grants/projects/researchers and add tags with commas. Confirm creation and invitations → Save writes a real Google event. Recurrence is not supported.

### 3. Recordings, transcripts and manual links

Google Meet recording/transcription availability depends on the host's Workspace permissions. LUNA does not start recording. When Google adds recordings or Gemini notes as event attachments or description links, the next sync displays them as meeting resources. Google processing may take time. If missing, first inspect the original event attachments. For a manual link: Drive file → Share → Copy link → LUNA event/material → Edit classification and related items → Recording or Transcript/notes link → Save. Viewers also need permission on the original Drive file.

### 4. Verify changes and avoid duplicates

After saving, inspect the original Google Calendar event for title, times, guests and Meet link. If LUNA asks you to review the result, check Google before trying again; uncertain requests are not retried automatically. Later Google description/attachment updates arrive on the next sync. LUNA research classification, local notes and manual links remain separate.

## 5. Choose LUNA Lab within your Google account

1. Calendar → Create event → Connect Google event creation permission for the intended connection.
2. Select the Google account and allow both event access and calendar-list reading. Reconnect once if previously authorized without this additional scope.
3. After authorization, choose LUNA Lab → Save. Only calendars the account can edit are listed. With existing authorization, use Choose calendar in this account on the creation page.
4. Check that Google calendar now shows LUNA Lab. Dates default to today, 09:00-10:00 Seoul time; adjust the times before saving.
5. Selecting a different calendar creates a separate connection and preserves existing events. Moving existing events between calendars is not provided.

## 6. Edit, delete and search tags

1. Open an event by its title. Edit event changes Google title/time/location/description and sends attendee updates after confirmation. Edit classification and related items changes LUNA grants/projects/researchers/tags and local notes.
2. Only platform administrators or verified Google calendar owners can delete. By default this hides the event in LUNA, including after later syncs, while preserving Google.
3. Deleting the Google original additionally requires the current user to have authorized the Google account and Google CalendarList to report owner access. Select Delete Google Calendar original too, confirm and Save; attendees receive cancellation. Platform administrator status or shared writer access alone is insufficient.
4. Older connections may lack an authorizing-user record. The owner must reconnect Google event creation permission once before ownership can be verified.
5. Saved tags appear as #tag at the top. Click one to search tagged events in Materials, or use Materials → Advanced search → Tags. Include Event if filtering by material type. Concurrent Google changes require reopening the form before retrying.

## Check invitation addresses and the latest meeting resources

1. Search names under Invite researchers and select email chips. Google Calendar invitation email in the profile takes priority; an empty value falls back to the profile email.
2. Google invitations send real email. Review recipients, calendar, dates and the invitation confirmation before saving. Researcher metadata alone does not send invitations.
3. If Google has the recording or smart notes but LUNA does not, check the connection's last collection and automatic collection setting. Collect now if needed and verify completion.

4. Edit event supports switching all-day and timed events. Recheck dates and Seoul time before saving. Reopen the form if the Google original changed concurrently.

## Search terms and relevance

1. Enter several terms such as imagery eeg. Exact phrases rank first, then all terms, curated related terms and partial matches. Supported bilingual terms include EEG/뇌파 and imagery/심상. This does not translate arbitrary terms or find every semantically similar document.

2. Use "imagery eeg" for a required phrase, imagery OR eeg for either term, and eeg -animal to exclude a term. Quoted terms are not expanded. Use up to 200 characters and 16 terms. Suggested queries apply only when selected.

3. Choose Relevance or Default order. Date, grant, visibility, tag and current access restrictions stay in effect. Calendar items stay on their dates; the list below can rank by relevance. Announcements search only the published edition in the current language.

4. Basic search makes no external AI calls and consumes no personal or lab AI allowance. It searches the fields available on each page, such as a material title and stored body, without fetching inaccessible external files. Citation search uses titles only and preserves the conversation readers' access restrictions.

## 28. Generate, schedule and review reports

Generate and schedule reports, then review their evidence.

Open this guide online

### 1. Before generation

1. Verify management permission for the grant or project. Platform administrators, PIs and lead researchers generate within their managed scope.
2. Inspect the latest collected bodies, source dates and researcher assignments in Materials. AI does not independently discover every uncollected external document.
3. An administrator checks AI enablement and runner readiness in Lab management → AI/Drive settings. Review your allowance under Profile → AI usage.
4. Enter plans, objectives and desired reporting perspective under Research goals/AI context. These are reference context, not model retraining.

### 2. Generate a report manually

1. Open Reports → New report → Draft with AI.
2. Type and select grants, projects and researchers from the dropdown. Add multiple items and remove with ×. Review the scope to avoid unrelated evidence.
3. Choose daily, weekly, monthly or a custom period and check the start/end dates when shown.
4. Submit, then follow Pending, Running, Ready or Failed in the report list. Avoid repeated submissions while waiting. The same target and period may reuse an existing report; choose a new version only when regeneration from updated evidence is intended.

### 3. Create or edit a schedule

1. Open Activity reports → Schedules → Create schedule.
2. Choose a lab and target. Recurring schedules currently accept one grant, project and researcher each. Use manual generation for a one-off combined target.
3. Choose the period and execution hour: 8 means 08:00 and 20 means 20:00, Seoul time. Daily covers yesterday; weekly runs Monday for the previous Monday-Sunday; monthly runs on day 1 for last month.
4. Check Enable this schedule and Save schedule. Verify Active and Next run on the card. Use Edit schedule on that card to change its settings.

## 4. Disable a schedule clearly

1. Choose Pause schedule to stop temporarily, or Cancel and delete to remove the schedule; no other fields need to be filled.
2. For pausing, verify Paused and No next run. Unchecking Enable this schedule in the editor and saving has the same effect.
3. Pending requests belonging to that schedule are cancelled. Running or completed reports are not stopped or deleted.
4. Paused schedules stay in the list with Resume and Edit controls. Cancel and delete removes the schedule itself while retaining existing reports.

## 5. Review, annotate or delete a result

1. Open a completed report and check its subtitle targets and reporting dates.
2. Follow numbered evidence links to expand source material. Distinguish dates, actual work and decisions from AI interpretation. Missing evidence does not establish that no work happened.
3. Edit metadata to save tags, notes and assigned researchers. Add multiple researchers by typing and selecting. Assignments do not recalculate the original generation scope.
4. Confirm deletion to move an unwanted report to Trash. Disabling a schedule and deleting a report are separate; Drive copies and operational backups are not removed.

## 6. Choose instructions, privacy and completion email

1. Select and load an instruction template, then add the questions for this report. Follow Reusable report instructions for details.
2. Select Only me for private review and the full-body email option if needed. Delivery uses the author's profile email.
3. Review multi-scope targets and dates for immediate generation. Schedules use their saved targets, time and instructions on later runs.
4. To write the body yourself, use Write report rather than AI generation. See the direct-writing/privacy/email chapter.

## 7. Read the common sections and handle failures

1. New ordinary AI reports use Summary → Progress against plans → GitHub development → Researchers/collaboration → Needs verification → Paper preparation → Proposed next steps. Magazine articles use an editorial format.

2. Open sources to distinguish observations, interpretation, proposals and undecided matters. Missing evidence may leave a section limited instead of generating invented content.
3. If a run failed, inspect the same report first. Do not submit again if a requested retry has already completed.
4. For persistent queues/failure, check collection, scope, period, AI enablement and allowance. Send Support the report URL and request time. A Notion publishing failure is separate from AI report generation.

## Write AI and manual magazine articles

1. At the bottom of the magazine home, use Write an AI article or Write an article manually. These compact links are available to lab, grant or project managers.
2. Select the lab, related grants/projects/researchers and evidence dates. Scoped managers must select a scope they manage. Multiple selections are supported.
3. AI: enter a topic and instructions describing the questions, process, readers and structure. Generate the article and refresh its status. Enabled AI, available quota and accessible collected evidence in the date range are required. Narrow the scope if too many sources are selected. Finished articles have a whole-article summary and source list.
4. Manual: enter the title, citation-free summary and Markdown body. Use heading, bold, table and preview controls, @mentions and [[document search. Save a draft for yourself, continue through Edit article and publish for the selected audience. Authors can edit their published articles.
5. Editorial articles are separate from researcher awards. Related researchers do not automatically receive badges or prizes. The archive distinguishes AI and manual articles.

## Find reports and reproducibility results

Switch between All reports and Reproducibility results in Reports. In New report, choose Write manually, Compose from results or Draft with AI. To reuse results, select sections on result pages, then inspect their provenance, pinned versions and order in Review selection. When editing an existing result report, use Add to this report and save to apply changes. Check the Notion destination for manual writing and the platform destination for result compositions in their respective editors.

Settings → Reproducibility connections manages your keys across projects. Project settings → Connections and research tools opens the same manager with that project selected. Participants without project-settings permission can use Settings or Reports to manage their own keys. A key belongs to its issuer and one project; revoking it does not delete existing results or reports. Opening a menu does not issue a key, run AI or send anything to Notion.

## 29. Reproducibility results, Python and reports

Manage project connection keys, register Python/Colab outputs and select sections for reports or Notion.

[Open this guide online](#)

### 0. Try an example before installing

Choose a project in Connect results and select Try a synthetic example to create two private pages from real CPU calculations. Inspect charts, metrics and tables, then combine sections into a report. Repeating this reuses the same project examples. It uses no research data, GPU, AI or Notion delivery.

### 1. Register outputs from existing code

In Reports → Reproducibility results → Connect results, choose the project and select Copy connection cell. Paste it into the first Colab/Jupyter cell or near the beginning of your Python file. It prepares SDK 0.3.0 from a pinned wheel/SHA-256 and opens browser connection approval. Check the account, project and confirmation code before approving. A valid connection stays in current-process memory for reuse across that project's results. Connecting does not send results; reconnect after restarting the runtime. Never put key secrets in source, output, Git or AI consultation. Colab must use the actual platform HTTPS address, not your PC localhost.

After your existing computation, add `import luna_repro as repro; selection = repro.select()`. It recommends supported local variables, tables and figures; some may be preselected, but nothing is sent automatically. Check titles, descriptions and order, preview, then choose Send reviewed. If you already know the objects, `repro.publish(fig, results_df, metrics, title="Experiment results")` explicitly sends them. Default recommendation uses no AI and does not judge scientific importance or experiment success.

In an ordinary .py file or without widgets, use `selection.summary()`, `choose()`, `option()`, `preview("preview.html")` and `publish()`. Preview freezes selected content so later variable changes cannot silently alter it. A final cell cannot recover old closed figures, vanished variables or computation provenance. Opt into `capture()` before a figure closes or retain a file. For precise sections, sources and execution evidence, keep `page = repro.page(project_uuid, "Experiment results")`, `page.section("quality").metrics(metrics)`, and `repro.publish(page)`. Manual connection keys, environment variables and existing ZIP installation remain in the folded connection settings.

### 2. Read results, code and verification status

Combine text, numeric metrics, tables and PNG/JPEG images freely. Use Full preview to inspect results in order and follow compact provenance links to code, expanding details when needed. Each section can carry code links and registration locations. Automatic references identify the logging file/line, not the computation origin. Supply the calculation function, pinned commit and notebook cell explicitly. Open execution environment follows your Colab or server link; the platform does not dispatch training, inference or GPU jobs.

Imported outputs, reported executions and synthetic examples are separate. Execution and evaluation completion are producer claims, not independent reproduction or researcher approval. Version comparison shows differences between matching metric names and does not declare an improvement. Check matching data, splits, alignment and evaluation conditions.

### 3. Compose reports, copy and publish to Notion

Select the sections you need on result pages. Open Review selection in the selection bar to inspect results and versions from multiple pages, adjust their order and create a report. When editing an existing result report, use Add to this report and confirm changes by saving. Platform compositions begin as owner-only drafts. Selected result versions remain fixed; later experiments do not silently rewrite a report. Manage title, description, grant/project, researchers and tags through the existing material view. Result access follows its execution project and audience; changing search metadata does not expand it.

Copy link copies the section URL; Copy content and source provides supported formatting and provenance. If the destination cannot paste images or formatting, use Publish to Notion. Review the destination and automatic preview in one panel, then publish. Creating a child page is the default; append is in Publication options. Images are uploaded to Notion and metrics/tables become native blocks. Notion audiences are independent; revoking LUNA access does not automatically delete external copies.

### 4. Continue in AI consultation or ChatGPT

Choose Review with AI on a section or download its review context and deliver it to your preferred conversation. Inspect the pinned version, metrics, explanation, source references and omitted evaluation together. Connected consultation rechecks current source access. Reading or downloading does not call AI; actual consultation follows the existing AI budget and permissions.

In ChatGPT or Codex, explicitly provide an accessible result link or the downloaded context and ask: “Check the code provenance and evaluation conditions, identify limitations and propose the next experiment.” Work only through tools actually exposed by that host. A link does not imply that all external code was read or private repository access was granted. Generated code, synthetic checks, real execution and research success remain distinct.

### 5. Cost, limits and recovery

Registering, reading, composing or sending results to Notion does not itself call AI or start a GPU. Server storage/transfer and existing provider charges are separate; zero cost is not guaranteed. Keep models and large datasets on existing Colab, servers or Drive and send only selected small outputs. SDK JSON is limited to 16 MiB, each image to 8 MiB, and tables to 2,000 rows/100 columns. Notion export accepts 40 sections, 10 images and 99 rows/40 columns per table.

On transport failure, the SDK preserves the selected bundle in `.luna-repro/outbox` and returns pending. Treat that folder as private research data. Fix the connection, key or permissions, then explicitly retry with `python -m luna_repro sync`. Reusing a request does not duplicate its result. For authorization errors, check project membership, token expiry and revocation; for HTTP 413, reduce selected content or image size.

The default project quota is 1 GiB of result JSON and image bytes. Retrying the same request consumes no additional storage. HTTP 409 storage errors require reviewing selected outputs or administrator storage policy; HTTP 429 requires waiting before retrying. Default intake limits are 120 requests per user and 600 per project per 15 minutes, separate from AI daily allowance. Earlier results are not automatically deleted to free space.

Notion delivery distinguishes queued, sending and confirmed outcomes. An uncertain write is not automatically repeated. Inspect the destination and connection before recovery; it is not reported as published until confirmed. Platform versions and bounded image bytes are included in existing encrypted operational database backups. Private results are not automatically copied into the lab readable archive.

## 6. Distinguish Pasa and synthetic examples

The bundled examples/repro directory separates a deterministic CPU synthetic check from a Pasa output-folder adapter. The adapter reads an existing run manifest, selected renders, camera diagnostics and any available evaluation table. It neither runs the original inference/training code nor fabricates omitted evaluation.

Runs without ground truth or with skipped evaluation remain not evaluated. Preserve failed camera diagnostics and do not relabel partial inference timing as whole-pipeline runtime. Actual Pasa reproduction requires separately accessible data/models and a CUDA environment. Synthetic checks, imported outputs, a new Pasa execution and independent verification require distinct evidence.

## 7. Account closure and research evidence

Withdrawal ends login and project access and revokes reproducibility connection keys. Results, immutable versions, images and authorship history are retained. Queued Notion delivery checks current access immediately before execution and stops when access has ended. Copies already sent to external services are not automatically deleted.

Accounts with reproducibility results, composed reports or Notion delivery history require preservation or transfer review before permanent erasure. Erasure does not rewrite the authors of immutable evidence. Accounts with connection keys but no research history can erase those keys through the existing permanent-erasure process.

## Find results and manage connections

Switch between All reports and Reproducibility results in Reports. Use New report for Write manually, Compose from results or Draft with AI. In Reproducibility results, Connect results opens Python/Colab guidance; Manage connections opens your issued keys. Opening a menu does not generate a result, key or AI report.

Settings → Reproducibility connections lists your connections across projects. Project settings → Connections and research tools → Manage connections opens the same manager with that project selected. Participants without project-settings permission can still manage their own connections from Settings or Reports. Project-management permission does not automatically permit viewing or revoking

another user's keys.

A key belongs to its issuer and one project and can be reused for multiple results and runs in that project. Use separate names for Colab and servers if useful; names do not authenticate devices. The secret is shown only when issued. To replace a lost or old key, issue a new connection key, update the runtime, then revoke the old connection. Revocation preserves results, reports and version history. Expired, revoked and current project access unavailable are distinct. You can revoke your own key in the central manager even after losing project access. Readers use their own login and current sharing permissions, not the author's key.

## Step-by-step practice, syntax and downloadable manuals

Detailed practical guide: <https://lunablab-ai.com/workspace/repro/guide/?lang=EN>

View the finished example without installing: <https://lunablab-ai.com/workspace/repro/demo/?lang=EN>

Follow Colab, Windows or Linux setup, the final-cell workflow, incremental sections, text/metrics/table/image/figure syntax, code cells/lines/pinned commits, versions, reports, Notion, AI review, the Pasa adapter and upload recovery. Download standalone HTML/Markdown manuals and runnable `walkthrough.py/walkthrough.ipynb` in the practice ZIP. Read the HTML offline or print it to PDF in your browser.

Inspect the public synthetic example first, then create your own private project demo in Connect results. Code practice saves locally by default; only explicitly choosing to publish uses the connection key. Distinguish a published receipt from a pending local bundle. Neither synthetic example measures research performance or independently verifies reproduction.

Public SDK installation and upgrades:

<https://github.com/lunablab-ai/luna-repro-sdk/blob/v0.3.0/docs/installation-en.md>

SDK 0.3.0, the `luna.repro.page.v1` contract and the platform version are independent. Public SDK code does not expose keys or research results; installation does not alter existing results, connections or history. Record the installed version and requirements before changing it; test synthetic examples in a new environment. To roll back, reinstall the recorded prior wheel and SHA-256 in that environment while preserving result JSON and outbox files. SDK/CPU practice needs no separate GPU or paid execution service; Colab resources, servers, storage, networking and optional AI have separate costs.

For ChatGPT or Codex questions, specify SDK `luna_repro 0.3.0`, the `luna.repro.page.v1` result-page contract, selected result JSON and error code to avoid confusion with legacy ReproKit events. Never include connection keys in prompts or attachments. Provide review context or an accessible section link, then continue within the actually exposed tools and current permissions.

## 8. Arrange narrative cells freely

In Compose from results, choose Add narrative cell. Use Markdown for the background, methods, data, analysis or limitations where needed. Move result and narrative cells with the up/down buttons or drag their handles on desktop. Add results can bring in sections from other reproducibility pages. Save your changes. This editor does not modify measured values inside result cards.

Saving records a report revision. Open Version history to inspect and restore earlier content; restoration creates a new revision. If another window has edited the report, review the current version first. Removing a result cell only removes it from this report; its source result remains.

## 9. Draft with AI from selected evidence

Save your report, then open AI draft. Select earlier materials or reports that you can currently access in the same lab, and add instructions. Images and saved source code are sent only when their respective options are selected. A link alone does not fetch external repository or Colab code. Review the selected counts, input scope and usage policy before requesting the draft. Reading and previewing do not run AI; generation uses existing AI access, daily allowance and request policy.

AI proposes a title and commentary on background, methods, data, experiments, analysis, limitations and discussion. Unsupported details must remain open questions. Select the proposed narrative cells and title you want to apply, then review them yourself. A changed report or evidence blocks overwriting with an outdated draft. Refine this cell with AI preserves its topic and appropriate length, updating only that draft cell with Undo. It never changes the title, other cells or result order; saving the report stays explicit. Edits made during generation require comparison before applying. Follow real progress and elapsed time for the same job. @mentions and [[material/report citations]] work in newly added cells; cited source text uses current access. This does not perform training, inference or independent verification.

## 10. Move results to Trash and restore

For a result you own, choose → Move to Trash on its card or detail page. Review the impact on referencing reports and ongoing Notion delivery before moving it. It disappears from normal lists, search and reads, while its ID, versions, images and reference history remain. Restore returns the same ID. Management privileges alone do not expose another person's private result.

An SDK retry cannot automatically restore a trashed result. On result\_trashed, explicitly restore your result in the platform before retrying. Ongoing Notion delivery requires inspection; external copies are not automatically erased. Referencing reports, AI and downloads recheck current source access after revocation or trashing. Permanent erasure of research evidence is not offered.

## 30. Write reports, control privacy and receive email

Write a report, attach files and control its audience.

[Open this guide online](#)

### 1. Choose AI generation or direct writing

1. Open Reports. Generate AI research report requests an evidence-based AI draft; Write report lets you enter the body yourself.
2. For AI generation, follow Generate, schedule and review reports. For direct writing, follow the steps below. Missing generation controls can indicate missing scope-management or AI permission.
3. Direct writing requires an eligible Notion connection you can manage and publish through. If the list is empty, ask the connection manager for access to a parent page in the same lab and content insertion permission.

### 2. Write and review a Markdown report

1. Open Reports → Write report. Enter a title and select the Notion connection, checking its name and scope.
2. Select related grants, projects and researchers; enter dates and tags. Leaving grant/project selections blank uses the connection's scope. Select the actual reporting audience because scope affects access.
3. Enter Markdown: **##** Weekly summary creates a heading, **-** Experiment setup creates a list item, and **\*\*Important\*\*** emphasizes text. Use **@username** or type **[[** and select a suggested document to insert mentions and citations.
4. Preview headings, tables, links and paragraphs, then save the draft. Draft storage and completed Notion publication are separate steps. Review the destination and content before publication.
5. Follow the Notion publishing chapter for the final action. After publication, continue editing the original in Notion.

### 3. Set or change Only me

1. Select Only me before generating an AI report or saving a direct report to restrict platform access to its author.
2. After creation, the author can turn Only me on or off in the report's visibility controls and save. Another administrator cannot change it on the author's behalf.
3. Turning it off restores ordinary scope-based access; it does not publish the report to the open internet.

4. In Reports or Materials advanced search, choose Only my private reports. Exclude private reports removes these from results.

5. This controls platform access. Previously sent email, downloads and Notion copies have separate access rules and cannot be recalled by this switch.

#### 4. Receive the complete AI report by email

1. Check that your profile email is an address you can access.

2. In AI report generation, select the option to email the full body after completion. This is a per-report choice, separate from ordinary notification email.

3. On completion, delivery is attempted to the author's profile address. Sender configuration, policy, sending limits and current access must permit it.

4. If missing, check report completion, your address, spam and sender configuration. Uncertain attempts are not automatically repeated. The completed report remains available on the platform regardless of email delivery.

#### Insert files and images inside the body

1. Open a new wiki page, Edit page, a new/edit proposal, Write report, or a manual magazine article.

2. Select the Notion connection for a report or the lab for an article first. That lab needs a configured private Google Drive archive.

3. Click the insertion point inside the body. Choose Attach file for PDF, DOCX, HWP, TXT, Markdown, ZIP and other supported files, or Add image for PNG, JPEG or WebP.

4. Limits are 20MB per file, 1.5MB per image and 20 attachments per document. Once Drive storage is confirmed, a download link or image Markdown appears at the chosen position. Images are normalized to safe JPEGs.

5. Check Preview, then save/publish or submit the proposal. Use Retry attachment or Cancel attachment if transfer fails. If Google archive authorization has expired, keep the editor open, use Open archive account settings in a new tab, reconnect the existing account and retry. No new folder is needed. Ask your lab administrator to reconnect if you cannot manage settings. Finish uploads before saving.

6. Before saving, only the uploader can access new files. Pending proposal files are available only to the author and reviewers. Saved files follow the document audience, including Only me and revoked memberships.

7. Notion exports and wiki ZIPs retain authenticated platform links for new attachments. They do not make files public or copy large originals into the ZIP. Download originals using their links. Read-only collected documents are still edited in their original service.

# 31. Reusable report instructions

Save and reuse instructions for recurring report tasks.

Open this guide online

## 1. Load a built-in template

1. Open Reports → Generate AI research report, or create/edit a report schedule.
2. Open the instruction template selector. Choose PI monitoring, Grant/project progress review or Researcher mentoring preparation.
3. Load it and read the inserted text. Preserve existing instructions before loading and review any replacement prompt.
4. Add your goals, experiments or mentoring questions. A template does not select the reporting scope or period; check those fields separately.
5. Submit generation or save the schedule. Instructions remain subject to the common report sections and evidence/access rules.

## 2. Save, edit and hide personal templates

1. Write the instructions, enter a descriptive name such as Weekly PI mentoring, and save as a personal template.
2. Select that name in a later generation or scheduling form to reuse it. The saved list belongs to your account.
3. In template management, edit the name or text and save. Names must be unique within your account; up to 50 templates are supported.
4. Turn off Use in selection list to hide a template. Re-enable it to use it again.
5. Editing a template does not update existing schedules. Open each relevant schedule, load the revised template and save the schedule. Completed reports remain unchanged.

## 32. Publish reports, materials and events to Notion

Preview a new Notion publication and handle uncertain results.

[Open this guide online](#)

### 1. Check the connection and permission

1. Open the report, material, document, data or event detail page. You need both access and the applicable management permission.
2. Find Notion publish near the bottom, usually beside Move to trash. If missing, check whether this source type and your role support publication.
3. An eligible connection in the same lab is required. Ask its manager to verify content insertion permission, an active connection and integration access to the parent page.
4. People who can access the Notion parent may read the published content. Platform Only me and project permissions are not copied into Notion.

### 2. Select destination, preview and publish

1. Choose Notion publish, then select the destination by connection name and scope.
2. Review the title, body, related context, original link and parent page. The operation creates a new page for this source.
3. Confirm only when content and destination are correct. Wait and check status while processing. It does not send file binaries or all linked documents.
4. On success, open the Notion link and check the title and body. Continue editing in Notion; the platform does not automatically overwrite both copies in both directions.
5. The confirmed Notion page ID prevents the published output from reappearing as a duplicate collected material. Existing documents are not removed merely because their titles match.

### 3. Handle failed or uncertain publication

1. Distinguish a confirmed failure from Could not confirm publication. An uncertain state may mean Notion created the page even though confirmation was lost.
2. First inspect the parent in Notion for a new page and use Refresh status in LUNA. Do not repeatedly create new reports or publish attempts without checking.
3. For a confirmed permission rejection, correct insertion capability, parent-page sharing and expired credentials. Retry only where the interface offers it for a definite failure.

4. If the source or connection changed during review, reopen the review screen and check the current content and destination.
5. For persistent uncertainty, send Support the source URL and time of action. Do not include secret tokens or the entire private source.

## 33. Print, PDF and document exports

Print or export an accessible page and check what is included.

[Open this guide online](#)

### 1. Open the export menu

Open a report, material, grant, project, researcher, event or conversation page. Select the download icon (Export) near the top. Share, Export and Favorites use matching icon buttons; hover to see each label. Choose PDF, Markdown (.md), or Text (.txt) to download. Configuration and credential forms do not have this menu.

### 2. Print a document

Select Export → Print view. Review the title, body, metadata and sources in the new tab, then select Print / Save as PDF. Choose a printer or Save as PDF in your browser. Navigation and editing forms are excluded.

### 3. What is included

Only document content you can currently view is included. Discussions, comments, likes, follow controls and editing controls are excluded from every export format. Search, calendar and list exports retain current filters and the current result page; other pages are not merged. Evidence includes visible previews and links. Export a material detail separately for its full displayed content. Linked attachments, images and external files are not embedded. Original content language is retained without translation or AI tokens. Export does not modify Drive archives.

## 34. Research messenger and AI consultation

Start conversations, understand read receipts and manage history.

Open this guide online

### Unread badges and receipts

1. Conversation, Research chat and AI consultation badges count currently accessible unread entries. Notification-inbox read status is separate.
2. Messages actually displayed in the focused conversation are marked read automatically. Viewing the directory or leaving older messages unloaded does not acknowledge them.
3. Open Read or the people icon beside a heart to see readers or likes. New or revised messages may need acknowledgement again.
4. Explicit Mark conversation read acknowledges its displayed boundary. Later arrivals remain unread.

### Open conversations and optional metadata

Close a conversation with × in the list or Close this conversation in its menu. This only hides it from your open list; it does not leave the room or delete history. Reopen it from /workspace/chat/. A new message or AI response brings it back. The top × hides the entire panel.

For research conversations, search and select participants. AI consultations start privately. Both use one optional input for multiple grants, projects, researchers and tags, or a typed topic. #tag input is supported. Research rooms include active participants of selected grants and projects and selected linked researchers at creation. Review the participant preview first. Topics and tags do not add people. AI metadata never invites other people.

AI uses sources matching both current read access and the selected context. Values of the same type are alternatives; different types must all match. Leave blank to search all readable sources. Missing evidence never broadens the scope, including through remembered advice. Existing history and scope are retained; start a new consultation for a different scope.

### Cite a page like a wiki

1. In the message box, type [[ and part of a title, such as [[Pasa.
2. Choose a grant, project, material, report, event or researcher using arrow keys and Enter, or click.
3. The input receives [[title|internal path]]. Add multiple citations and send to display clickable titles.
4. Research chat suggests only targets accessible to every current participant. AI consultation uses your own access. Search is by title.

5. Add tags with comma or Enter and remove with ×. Metadata for a new conversation is optional and separate from the currently selected chat.

## AI errors and researcher activity reviews

1. For a failed consultation, inspect the explanation in Full history. Resolve permission, allowance or readiness issues before Retry. Retries also count as attempts.

2. Researcher reviews are scheduled according to the configured evidence window and refresh interval. Repeated profile reloads do not run them.

3. Profile summaries and detailed reviews have different audiences. The subject and platform administrators can read details; evidence links require current access.

4. See AI assessment for role-specific criteria, evidence periods and unassessed items. Reviews support mentoring and do not prove experimental success or objective personnel ratings.

## Start: create channels and groups

1. Open the floating conversation button → Research chat → +. Optionally name it and select people, grants or projects.

2. Selecting a grant/project can include its current participants. Review the actual audience before creating. Topics and tags alone do not invite people.

3. Choose Start. A researcher's profile also offers a direct-conversation button.

4. Private groups are readable only by current participants. Nonparticipant platform administrators do not have blanket group access.

5. Close only folds the panel. Leave conversation ends participation and subsequent access and notifications.

## Manage, search and download history

1. Open Profile → Settings → Conversation management or Full history in the panel. Search names/content and filter by type, dates and context.

2. Open chat opens the side panel; Conversation settings manages the name, participants and automatic summary. More separates download, clear, leave and trash actions.

3. TXT/JSON exports include up to 2,500 entries. Split larger periods and download attachments through their links.

4. Only you manage your private AI consultation. Group administration still requires current participation and applicable authority; nonparticipants cannot bypass privacy.

5. Clearing, trashing and leaving have different effects. Review the confirmation. They do not jointly erase other people's preserved messages, existing reports or external copies.

## AI titles, commands and read-only sharing

1. Authorized AI users click + in AI consultation and optionally select context metadata. The title is optional: leaving it blank uses the beginning of your first question. Rename it in Name and settings.
2. Type / to see commands. Typing /report w narrows the list. Select with arrow keys and Enter or the mouse. Autocomplete does not send anything; review before sending.
3. Use /report daily, /report weekly, /report monthly, or two dates. Replace YYYY-MM-DD placeholders with real dates. Report permissions and usage limits still apply. In a new private consultation, use the Reports page to choose generation conditions. Existing grant/project consultation commands remain available.
4. AI conversations are private by default. Enable Share read-only with channel participants in settings to share the history link with that channel. Readers cannot ask, edit or delete, and answers are hidden when they lack evidence access. Disable the option to revoke sharing.
5. Use the small share icon on the history page to copy its link or share to another channel/app. A link never grants access. Current research channel messages are searchable AI evidence; private AI answers are not independent proof of research progress.

## 1. Chat while browsing

1. Click the floating conversation icon at the bottom right to open the right panel. Toggle it again or press × to close. On mobile it uses the full screen.
2. Open a conversation from the list or press + to start one. Grant/project rooms are not automatically created. Navigation restores the conversation selection and typed draft in the same browser tab. Unsent file selections must be reselected after navigation.
3. Type and send a message. Type @ followed by the start of a username to select a participant. Mentions and replies create notifications for the recipients.
4. The open panel refreshes about every five seconds without forcing you to the bottom while reading older messages. Choose View history for full-channel search, editing and exports.
5. Grant participation does not automatically grant access to every child-project channel. Obtain approval for the relevant project.

## 2. Replies, search, citations and sharing

1. Choose Reply / details to open a stable message URL. Replies notify the original author inside the platform, and mentions notify selected participants. Email requires administrator Gmail setup and recipient opt-in. SMS is not provided.

2. Enter words in channel search. Expand advanced search for author, tags and date boundaries.
3. Copy a citation link and paste it into another conversation. On a report, material or event, choose the small share icon → Share to channel → target channel. Review the prefilled link and message before sending.
4. Recipients must independently have access to the linked page. Sharing does not fetch or preview unauthorized source content.

### 3. Edits, deletion and confirmed actions

1. Open Edit / delete on your own message, or on a message in a scope you manage. Save the edit or delete the message. Refresh if someone changed it concurrently.
2. Deletion leaves a tombstone; earlier versions are retained. Editing clears the previous decision confirmation. Authors and managers can inspect revisions.
3. Managers can mark a confirmed decision, confirmed action or unresolved question. Discussion does not establish experimental success.
4. A confirmed project action offers links to connected GitHub boards. Review the proposed issue and explicitly confirm the remote change. A stable request identifier prevents repeated creation from the same message. If a network outcome is uncertain, check the remote issue first.

### 4. Grant AI conversation access

1. A platform administrator opens Profile → Lab management → AI access and daily limits. Being a grant or project manager does not grant authority to enable AI access for others.
2. Select a user, enable AI conversation access, choose allowed models and maximum reasoning effort, then save. Choose from the available models and limits shown in the form. Use Change on an existing permission card to edit it.
3. To revoke, clear Enable and save. Queued or running requests are invalidated. Owners can still read their completed private answers while they retain conversation access. New requests and shared answers require current permissions.
4. The daily AI allowance is shared by reports, translations and consultation. Enabling chat does not expand data access or report management rights. Actual model availability depends on the server Codex account.

### 5. Consult AI and request reports

1. Authorized users choose Conversations → AI consultation → +, optionally add a title and context, then save. Private consultations are excluded from shared lab Drive archives.
2. Ask for confirmed progress versus unresolved issues, a researcher's evidence-backed activity, paper experiments or future research directions. Expand Model and effort to change settings for each question.

3. Wait for processing and use the new-response notice to refresh. Expand citations and compare the evidence. Retrieval searches readable materials, messages, events, reports and your own consultations for relevant excerpts; it does not claim to read all originals. Collect external materials before asking about them.

4. If you manage the scope, expand Generate research report in this conversation and choose daily, weekly, monthly or a custom range. You can also type `/report daily`, `/report weekly`, `/report monthly` or `/report 2026-09-01 2026-09-13`. The report link appears in the conversation and in Reports. Equivalent requests reuse existing reports.

5. Ordinary natural-language consultation does not change settings, send messages or execute experiments. Use the explicit report command or control for creation. Your completed answers remain in history; shared answers, source links and new requests follow current access. Do not treat an old answer as current evidence; ask again when needed.

## 6. Files, exports and clearing history

1. Research chat and AI consultation support file selection, pasted images and dropped files: up to three files, 1.5 MB combined. Review the list before sending; see Compose and attachments.

2. Channel files require current participation; private consultation files follow that consultation's audience. Links grant no access.

3. Export history by dates, author and keywords as TXT/JSON. Download original files through their attachment links.

4. Clear/delete follows the confirmation scope. Deleted messages are excluded from new AI retrieval; existing reports and historical copies are separate.

5. Private consultations, direct messages and participant-only summaries are not all copied to the shared lab Drive. Operational DB recovery backups are distinct from app sharing.

6. Direct messaging is available; voice/video calling is not. Use a separate Google Meet event for meetings.

## Mentions, replies and daily limits

1. In research chat, type `@username`, choose a participant and send. Participants who can currently read the channel receive notifications. An explicit self-mention also creates a notification; ordinary messages you send do not notify yourself.

2. Click `@username` to open the profile, or ↗ next to a reply to open its original message. Code blocks and email addresses do not count as mentions.

3. In private AI consultation, `@username` is a reference only. It does not share the conversation or notify that person. Use research chat for shared communication.

4. Platform administrators: Profile → Lab management → AI access and daily limits → Change on a user card. Set AI conversation access, daily attempts, allowed models and maximum effort, then Save. Selecting another user reloads their saved settings.

5. Reports, AI consultation, AI translation and normal material analysis share the daily attempt allowance, including failures. On-device translation uses none of this allowance. Blank uses the role default; 0 disables usage. The counter resets at midnight in Seoul. Token-based limits are not provided. Enabling AI does not grant additional grant or project access.

## Free message translation, originals and reuse across devices

1. Use A↔가 below a message to translate between Korean and English. Switch to Original or Both to compare. Links, code, mentions and attachments remain intact.

2. Choose interface language, Korean, English or Always original in translation preferences. Your preference and private copies sync across devices. Reader translations remain private; only author-approved shared translations appear for other current participants.

3. Expand Translation preview below the composer, translate for free, review and opt to share with the original. Optional previews while typing never block sending the original. Recheck after editing.

4. Desktop Chrome uses the Translator API without consuming AI allowance. First use may require clicking the icon to download a language pack. Unsupported browsers including mobile display saved copies. Create a new copy on desktop or use Chrome menu → Translate yourself. The website cannot trigger or read Chrome menu translation.

## Read recent conversation and earlier history

1. A conversation opens on recent history. Expand earlier conversation or Load older messages to read more.

2. While browsing older messages, Show latest returns to new arrivals. Full-history search also searches folded content.

3. Automatic receipts cover displayed messages; Mark conversation read has a wider explicit boundary. Unseen older messages are not automatically acknowledged.

4. A message link opens its surrounding history after an access check. If unavailable, check current channel participation.

## Find unread messages in one place

1. Use the conversation badge or unread list to choose a conversation. Research messages and new answers in your AI consultations are distinguished.

2. Search and paginate the full unread directory for more; ordinary conversation lists also prioritize unread activity.

3. Displayed messages in a focused panel receive automatic receipts. Closing, listing or receiving push alone does not mark unseen messages read.
4. Network conditions may delay refresh. Notification read state, conversation read state and pending approval counts are separate indicators.

## Direct messages, lab channels and prior consultation

1. Start a direct message from a profile or choose a person in the conversation picker. Continue an existing conversation when one opens.
2. Use an existing lab channel or invite the intended participants to a new conversation. Check the channel and participants before sending something private.
3. AI consultation can retrieve accessible collected documents, GitHub Projects work, conversations and your own completed consultations. A connection does not mean unlimited reading of everything it contains.
4. Name the project, period and document and cite relevant sources. Open the answer's references to verify coverage and omissions.
5. Other people's private consultations and support inquiries are not research conversation evidence. Completed visual analysis of a PDF/image can also supply supporting notes.
6. Click the copy icon below an AI response. After Copied appears, paste the answer into your document.
7. If automatic copying is blocked, copy the selected response manually. Answer copying excludes the question and separate source list.
8. The history page copy icon copies both the question and answer.

## Manage participants and periodic AI summaries

1. Check current participants in Conversation settings. The creator can remove other participants, ending their access and alerts. Leave ends your own participation.
2. Review automatically included scope participants when creating the conversation. Tags do not automatically rejoin someone who left.
3. Enabled summaries become eligible roughly every 30 minutes when new messages exist. A current participant's AI permission, lab settings and allowance are required; completion is not immediate.
4. Open AI conversation analysis to inspect status and the accumulated summary. Pausing retains valid results. Edited or deleted originals may require reanalysis.
5. Summaries are participant-only AI interpretations, not confirmed decisions or research results. The summary task does not automatically read the full originals behind attachments or citation links.

## Compact message translation and deletion

Choose one translation language in the conversation's menu → Translation settings.

Each message translation icon follows this preference.

View original is a small link; Show with original and Share reviewed translation are inside .

Translation, editing and deletion share one row.

Message authors or channel creators with current management permission can confirm deletion with ×.

The composer translation icon opens a draft preview.

Send a reviewed translation alongside the original, or apply a plain-text translation to the input.

For links, code and mentions, send the original with its translation to preserve formatting.

Only free browser translation and saved copies are used.

Mobile browsers without Translator API can display saved copies; use Chrome menu → Translate for new translations.

No AI allowance is used.

## Zoom, pan and download in the image viewer

1. Select an image thumbnail to open the viewer. Use zoom buttons or the wheel on desktop and pinch gestures on mobile.
2. Drag an enlarged image to pan. For multiple images in one message, use previous/next. Close or press Escape on desktop to return to the conversation.
3. Choose Download separately in the viewer. It downloads the stored image through its authorized path and does not grant access to private attachments.
4. On small screens use controls within the safe area. Zooming cannot recover detail absent from an older low-resolution stored image.

# 35. Compose, cite and delete messages

Compose messages with citations, pasted images and dropped files.

Open this guide online

## 1. Compact composer and participants

1. Check Participants on the channel page. Grant/project channels use scope membership; groups use their selected members.
2. Write in the compact composer. Expand Attachments, links and tags for extra options. Preview supports headings (#), bold (\*\*), lists (-), inline/fenced code and links. Select Edit to return.
3. Use × to delete your own message or one you are allowed to manage. Channel settings → Conversation history offers clear and delete. Deleted content is excluded from the current view and AI search; historical archive versions remain.

## 2. Citations, mentions and AI context

1. Type [] and part of a title, then select a result. Research chat searches sources readable by every participant; private consultation uses your access.
2. Type @ and a name or username. A private AI mention is a reference, not a message or notification to that person.
3. Grants, projects, researchers and tags are optional context for new private consultations. Narrow them when useful and state dates and document titles in your question.
4. Existing scoped consultations retain their saved scope. Switching research channels does not automatically change an existing AI consultation's context.

## Insert files and images inside the body

1. Open a new wiki page, Edit page, a new/edit proposal, Write report, or a manual magazine article.
2. Select the Notion connection for a report or the lab for an article first. That lab needs a configured private Google Drive archive.
3. Click the insertion point inside the body. Choose Attach file for PDF, DOCX, HWP, TXT, Markdown, ZIP and other supported files, or Add image for PNG, JPEG or WebP.
4. Limits are 20MB per file, 1.5MB per image and 20 attachments per document. Once Drive storage is confirmed, a download link or image Markdown appears at the chosen position. Images are normalized to safe JPEGs.

5. Check Preview, then save/publish or submit the proposal. Use Retry attachment or Cancel attachment if transfer fails. If Google archive authorization has expired, keep the editor open, use Open archive account settings in a new tab, reconnect the existing account and retry. No new folder is needed. Ask your lab administrator to reconnect if you cannot manage settings. Finish uploads before saving.

6. Before saving, only the uploader can access new files. Pending proposal files are available only to the author and reviewers. Saved files follow the document audience, including Only me and revoked memberships.

7. Notion exports and wiki ZIPs retain authenticated platform links for new attachments. They do not make files public or copy large originals into the ZIP. Download originals using their links. Read-only collected documents are still edited in their original service.

## Paste images and drop files

1. Choose files or drop them into Research chat or AI consultation. Paste a copied image with Ctrl+V. Ordinary text paste still works.

2. Review filenames, sizes and image previews; remove unwanted files. The limit is three files and 1.5 MB total. Register larger files in Materials and cite them.

3. Attachments upload when you send the message or question. After failure, review the retained queue before retrying. Navigation or switching conversations may require selecting files again.

4. Resizing and text extraction have limits. Identify the important table, figure or page and compare the answer with the original. Attachment content does not authorize code execution or browsing its links.

## 36. Personal AI memory and source coverage

Use personal memory and understand source and analysis coverage.

[Open this guide online](#)

### Manage personal research memory

Use Personal AI memory to add, edit or forget preferences, decisions, hypotheses, research lessons and task-specific conditions. Project or grant scope remains private to you; it does not publish team knowledge. Current instructions override old preferences. Source-backed memories are not reused after their source changes, is deleted or becomes inaccessible.

Recall uses relevant excerpts from completed personal consultations and external discussions actually delivered through supported tools. Personal consultations, new research strategies, code consultations and daily plans/reviews share this service. Shared conversations, public reports and translation do not receive wholesale private history. Use the existing reviewed report sharing copy and recipient selection for team knowledge. Code knowledge retains the repository commit, source reports and verification status.

Excluding a conversation stops implicit recall elsewhere but preserves original history and saved outputs. Re-inclusion never overrides a previous reset cutoff. Memory reset and source deletion are separate. Previous outputs remain history, but forgotten memories cannot become future evidence.

Registration, search and exclusions use no extra AI calls. Context tokens are included in normal response usage. Selected memory text is capped at 8000 characters; this is not exhaustive recall. Source data, indexes, settings and exclusions belong in encrypted external recovery backups. Apply the latest exclusion ledger before reopening a restored application. Backup and provider retention periods are distinct from immediately stopping reuse.

### 1. Recall across conversation windows

1. Open Profile menu → Settings → Personal AI memory. Use personal memory across conversations is enabled by default.
2. Ask a question and wait for a completed answer. Your question and answer are stored for your account. In another window, mention the earlier topic or distinctive terms, for example: “Use the validation sequence we discussed to plan the next experiment.”
3. Existing completed consultations are indexed gradually. Unindexed history is still searchable and older matches have no fixed date cutoff.
4. Open references to check which conversations and materials informed the answer. This retrieves related history; it does not retrain the model or read every conversation in full on every request. Repeat critical requirements in your question.

5. Automatic memory uses your own completed consultations. Other users' private consultations and support tickets are excluded. Deleted conversations and currently inaccessible evidence are not reused. AI suggestions are not confirmed research facts.

## 2. Pause or reset memory

1. Clear the checkbox and select Save memory settings to pause recall, including implicit history within the same window. Re-enabling recalls history after the most recent reset.
2. Select Reset personal memory and confirm to delete saved memory and establish a new cutoff. Answers already in progress cannot restore pre-reset memory.
3. Original consultation history stays available for viewing and export, but pre-reset history is not automatically reindexed. New consultations establish new memory.
4. To delete the original conversation, use its history-management controls. Reset cannot recall previously downloaded or externally archived copies. Personal memory is excluded from shared lab Drive archives.

## 3. Make sources available for consultation

1. Register and authorize GitHub, Notion, Drive and other connections on the grant/project, then check collection status. Consultation searches currently readable collected materials, GitHub Projects work, reports, documents, events, channel/direct messages, notices, discussions and grant/project descriptions.
2. To collect an entire Drive account, select All Drive accessible to this Google account and set the lookback period to 0. Select a folder for narrower coverage. Scope participants may see collected materials, so consider personal files before selecting the whole account.
3. Save, authorize Google and select Collect now. Large listings continue over multiple runs. Inspect extracted content, source dates and omissions. An uncollected linked file has not been read by AI.
4. Include document names, scope, dates and key terms, or insert a [[ citation. All allowed sources are candidates, but a response receives selected excerpts with at most 80 evidence items and a byte limit. It does not read every file in full at once.

## 4. Automatic text and visual analysis, with limits

1. Edit the Drive connection and select Automatic analysis (default) under PDF, image and embedded document image analysis. Existing connections are eligible unless explicitly disabled. On request only stops automatic queue enrollment.
2. Text: TXT, MD, CSV, JSON and related text, DOCX, PPTX slides and notes, XLSX sheets and cached values, HWP 5 body text, HWPX, PDF text layers, ODT, ODS and ODP. Google Docs, Sheets and Slides are exported to Office formats. Macros and formulas never execute.
3. Visual analysis processes PDF pages, PNG/JPEG/WebP/TIFF/BMP files and embedded raster images in Office, HWPX and ODF packages. The connection owner needs AI permission, lab AI enablement and daily

allowance. Exhausted allowance retries the next day. Previously collected files are enrolled gradually.

4. Open File visual analysis on a material to inspect progress and notes. A unit is a PDF page or an embedded Office image. This does not fully render Office layouts, vector diagrams/SmartArt, HWP images or transcribe audio/video. Verify numbers and plots against originals.

5. Limits include 20 MB per file, 200,000 extracted characters, 300 PDF/image units, and pixel/compression limits. Google export size limits also apply. Encrypted/distribution HWP, HWP 3, legacy DOC/XLS/PPT and unsupported files may be unreadable. Export a readable PDF, DOCX or HWPX and reconnect it.

6. Analysis notes are distinguished from originals in consultation retrieval. Pending, failed or partial analysis is not complete coverage. Recheck collection and analysis after source or permission changes.

## 37. Create and download AI files

Create PDF, Word, CSV and Markdown files in AI consultation.

Open this guide online

### Choose a result format and download

1. Expand Result format beside the attachment control in AI consultation. Choose Text, PDF, Word, CSV or Markdown; the menu folds and shows the selection.
2. Describe the purpose, audience, structure and evidence, then send. For example: create a Word meeting-preparation memo from the cited sources.
3. Download the generated file below the completed answer. Pending, running and failed requests are not completed downloads. Check the reason and allowance before retrying.
4. Limits are one file per answer, 8 MB per file and 40,000 generated characters. Complex layouts and perfect formatting are not guaranteed. Review citations, numbers and tables.

### Input attachments versus generated files

1. Attached inputs are evidence; files produced through Result format are AI outputs. They do not overwrite the input originals.
2. Generated files follow the consultation's audience. Private consultations are private by default, and shared readers still require current access.
3. Generated files are not automatically copied to lab Drive, and ChatGPT files are not automatically saved in LUNA. Download what you need and use authorized material registration.
4. LUNA does not offer separate paid-API image generation. ChatGPT image/file tools are separate and depend on that account's features and limits.

## 38. Connect ChatGPT and execute actions

Read authorized sources and, with additional consent, prepare records and actions.

[Open this guide online](#)

### Quick start and examples

Private journal PC tasks — [Preview](#) → [approve](#) → [execute](#)

Check private journal intake on my registered default PC. Preview exact intake, transcription, draft or reviewed-publication effects first. Keep source bodies off the research server.

Daily research planning and check-in — [Preview](#) → [approve](#) → [execute](#)

Read my daily research briefing and review priorities and blockers. Show the content before saving or sharing.

Find materials and events — [Read now](#)

Find imagery EEG materials and summarize them with source links.

Check GitHub work — [Read now](#)

Show accessible open GitHub tasks and their last observed time.

Create a GitHub Issue — [Preview](#) → [approve](#) → [execute](#)

Show repositories where I can create an Issue, ask for its title and body, then preview it.

Schedule an event — [Preview](#) → [approve](#) → [execute](#)

Show available calendars, ask for the date, time and attendees, then preview the event.

Generate a report — [Preview](#) → [approve](#) → [execute](#)

Prepare a report for the last seven days, private to me, with no email. Ask for the grant and show a preview first.

Generate a private article — [Preview](#) → [approve](#) → [execute](#)

Prepare a research article draft. Ask for the lab, period and topic, then preview it without publishing.

Check generation progress — [Read now](#)

Check my recent generation jobs and whether their bodies are ready. Explain recovery for failed jobs.

Help and troubleshooting — [Read now](#)

Find a solution in LUNA help. If unresolved, prepare a support draft for my review.

Draft a bilingual notice — Review and submit on website

Prepare a bilingual notice draft and a web review link. Do not publish it.

Draft a support request — Review and submit on website

Draft a support request summarizing the problem and attempted fixes. Exclude secrets and research content; I will review and submit it.

Cloud and PC execution support and status — Read now

Show online operations and current registered-PC readiness for ANNA. Read execution using its existing ID; distinguish host exposure from actual completion.

Revise, build and privately save courses — Preview → approve → execute

Find my registered course and session, then preview these exact file edits. Preserve existing approvals, review term-wide private save effects, and verify the execution receipt.

Continue and save research work — Preview → approve → execute

Continue the current research work and save this discussion and delivered files to that same work. Identify a separate intent before creating new work.

Draft event and task wording — Preview → approve → execute

Prepare a draft for the next meeting in this research space. Show the target and cost first; do not register an event or send invitations.

Keep existing work — Preview → approve → execute

Bookmark the research work I am viewing in My saved items.

Research goals, literature and dossiers — Preview → approve → execute

Prepare a paper goal from the selected project and check attachment readiness. Show cost and changes before execution.

Code work, ablations and validation — Preview → approve → execute

Prepare a separate code workspace from this idea and show its validation plan and diff. Request approval before publishing.

Comments, chat, discussions and actions — Preview → approve → execute

Prepare a question for this material discussion. I will approve after reviewing the destination and body.

Action results and follow-ups — Preview → approve → execute

Read my action history and show responses to review. Show follow-up changes before saving.

## Create directly from ChatGPT

1. Refresh your LUNA connection and select it in a new ChatGPT conversation. Ask to create an Issue, Google event, report or private article. First use requires additional OAuth consent for that action; existing connections never gain permissions silently.
2. `action_options` lists actual repositories, boards, calendars, labs and input choices. `prepare_action` shows the destination, content, assignees, attendee emails, sharing and sending effects. Previews expire after ten minutes. Review the exact details before approving execution.
3. `create_github_issue`, `create_calendar_event`, `generate_report` and `generate_article` use the existing creation workflows. Invitations and GitHub notifications may be sent. Reports default to private; email delivery is opt-in. Articles remain private; public publication is separate.
4. Use the returned `action_id` with `get_action_status` for receipts and current report state. `queued` means accepted; `ready` means complete. After a timeout or `needs_review`, inspect status and the original before preparing another action.
5. Search and preview do not consume LUNA AI allowance. Actual report/article generation uses the existing personal allowance. Revoking a connection blocks future execution but does not undo completed actions.

## Connect and ask your first question

1. Open Profile → Settings → ChatGPT connection and AI output settings for the enabled address and walkthrough.
2. In ChatGPT developer mode, add a Plugins connection using <https://lunalaab-ai.com/mcp>, OAuth and DCR when offered. Availability depends on your account/workspace.
3. Review the LUNA account and read consent yourself. Do not enter API keys or DB passwords.
4. Select LUNA in a new ChatGPT conversation and search a title you can open in LUNA. Check search/fetch activity and actual source links.

Official connection guide: <https://developers.openai.com/plugins/deploy/connect-chatgpt>

## Read scope, disconnecting and costs

1. The basic connection searches and reads material accessible to your account. Additional consent enables Issue, event, report and private article creation. It does not delete existing records or dump the database. Current AI permission and project/channel access apply.

2. Retrieved content goes to ChatGPT. Disconnecting blocks future retrieval, not copies already transmitted.
3. Connections last up to 30 days. Check settings and reconnect after expiry or revocation. For missing results, check access and collection in LUNA first.
4. This read connection needs no separate OpenAI API key or API billing. ChatGPT's own subscription and usage limits apply. Its generated files are not saved automatically in LUNA.

## Expanded search, help and action preparation

1. Refresh tools in your existing ChatGPT connection and start a new chat. `search_materials` searches material/report bodies and event descriptions; `search_help/read_help` provide official bilingual guidance. `search_work_items` reads collected GitHub status, and `list_my_generation_jobs` lists your report/article jobs.
2. `list_workflows/prepare_workflow` prepare calendar, issue, report, article, bilingual notice and support forms. Review links work for your account for 30 minutes. Open the existing LUNA form, check the target, content and invitation/publication options, then submit. Preparing a link does not create, save, publish or notify.
3. In LUNA AI consultation, use Prepare from input under Extra features after describing an action in the message box. Search and preparation use no LUNA generation allowance; actual generation follows its feature policy.
4. Complete authentication in the provider screen. ChatGPT history is not automatically copied to LUNA. If help cannot resolve your issue, review and submit a support draft yourself.

## A report link exists but the body is empty

A request URL is not a finished body. Check `content_available`, `phase`, `completed_batches/total_batches`, `code` and `message` with `get_action_status`. Large evidence sets are analyzed in batches, then combined. Automatic batches use one personal allowance unit per request; an explicit retry uses another. Retry a legacy limited/split\_required job from its report page to resume the same report. Narrow requests above 500 sources to a project or researcher. Search follows current read access, report creation requires management permission, and platform administrators have broader access.

## Turn manually written plans into actions and results

1. Read manually written daily items, citations and mentions and request available action candidates. The web and ChatGPT share item identifiers/revisions; this does not expose another person's private notes.
2. For example: Read my second plan item, show comment/message alternatives, then draft the first option and ask before sending. Verify the actual destination and evidence.
3. Basic matching and saving text composed by ChatGPT do not make an additional LUNA AI call. Explicit LUNA recommendations/wording follow that feature's personal usage policy.

4. External writes require additional research-action consent and approval of the preview's exact destination/body. Check tools and permission on older connections and start a new conversation.
5. Results, linked replies and personal outcome records are shared with the web. They become evidence for newly generated reviews without converting uncertainty into completion. Continue on the web if a required tool is absent.
6. Detailed guides: <https://lunablab-ai.com/workspace/help/research-actions/?lang=EN> and <https://lunablab-ai.com/workspace/help/daily-research/?lang=EN>

## Collaborate from a shared report

Use Share and collaborate on a strategy report to review content, recipients and read or follow-up permissions. Report ID and immutable edition ID/version preserve lineage. Original attachments, private chats and GitHub access are not automatically shared. Open Shared research reports, then start your own discussion or code workspace. Share a saved result with the team by selecting an exact revision/checkpoint/code result and files to return to the author. Seen, review and adoption are separate human actions. Sharing, reading and exporting make no AI call; new generation uses the initiating user's allowance. Stable request IDs prevent duplicates. Revocation, expiry, withdrawal and project access changes affect subsequent access and continuation. Exported copies cannot be recalled. Notion queues a fixed edition under an authorized parent; unknown outcomes never auto-retry. Reconcile the existing page by ID. External researchers must first complete existing account and lab/project participation approval; member rollout is not automatically enabled.

## Research work links and graph

Use `consultant_thread_prepare(operation=save, values={space:{operation:read/graph/work_search/links/link_history,values:{space_id,...}}})` for the same canonical work. `work_search` accepts query and `kind=work/all_work/strategy/record`. Connect uses source and target as `kind:ID`, `kind=derived/reference/requires/revision/produces/validates` and an optional selector for the section or method. Add supports multiple `data.relations` entries containing source, kind and selector in one atomic save. Preview with the current revision and `stable_request_key`, then commit the exact review and `preview_token`. Include `link` and `link_revision` for explicit refresh or disconnect. Prerequisites are not automatic execution instructions. Reuse the same thread for a continuing task; only delivered answers/files are preserved. Server support and host tool availability are separate.

## Search action destinations

In Create action, choose a method and type a destination name, title or issue number. Personal AI and ChatGPT use `research_action_options` with `lab`, `kind` and `q`, omitting `target`, to search the same destinations. Results contain six items; when `has_more` is true, pass `next_offset` as `offset`. An explicit `kind` and `target` still reads destination context. Search consumes no LUNA AI allowance and never posts or sends. Saving a draft and approving the execution preview are separate. Actual ChatGPT use requires host availability of the tool.

## Private memory with cached connector tools

If personal account tools are not exposed yet, use `consultant_thread_prepare(operation=save, values={memory:{operation:read,query?,offset?}})` to read your memory and `settings_version`. Review mutations with `values={memory:{operation:prepare,settings_version,changes:{operation:save/forget/exclude/include/configure/reset,...}}}`, then pass the exact returned `preview_token` and `review` to `consultant_thread_commit`. Reads need no commit; changes require the current write scope. No new thread or paid AI call is required.

## Research with private memory

Ask “What do you remember about me?”, “Apply this preference only to this project”, or “Exclude this conversation from future recall.” Read memory through `consultant_account_read`, then review exact changes with `consultant_account_prepare(section=memory)` and `commit`. Use `memory_query` for retrieval and `memory_offset` for pagination. If the host has not exposed new arguments, continue on the Personal AI memory page. `consultant_thread_read` includes currently accessible personal context for the same work. Memory edits and retrieval invoke no additional model; conversation reasoning has its usual cost. LUNA cannot capture an entire ChatGPT account: only content delivered by supported storage tools is available. Private memories and conversations are never automatically shared with a team.

## Preserve multi-turn work in one place

Choose automatic preservation or sync on request when linking a discussion. Reuse the same LUNA thread ID across prompts. Native answers follow that policy; Sync available discussion preserves up to 50 accessible unsynced native answers per request. External ChatGPT/Codex must actually deliver relevant answers/files through tools. Read sync status with `consultant_thread_read`, then use `consultant_thread_prepare(operation=sync)` and `commit` for `configure/begin/chunk/finish/flush/confirm`. Cached hosts can use `operation=save, values={sync:{operation,values}}`. Scoped automatic preservation consent covers intermediate saves, subject to host permission prompts. It never authorizes confirmation, publication, GitHub writes or paid generation. The same discussion page shows current and confirmed results and cumulative history. Transfer files in 48KB chunks, up to 8MB, with 120KB conversation text and stable retry keys. Claim saved only after the ready receipt. Unseen messages and expired files cannot be recovered automatically. Saving and version management require no extra AI calls.

## Using the full feature set

Open Profile → Settings → ChatGPT connection and connect your own LUNA account. Select LUNA in ChatGPT and ask “Explain the LUNA features available to me.” Features either read information, execute after approval, or continue on the website. This guide covers server support; actual use also requires current permissions, external connections and host tool discovery. A missing tool is not a completed action. Continue on the corresponding platform page. Official connection guide:  
<https://developers.openai.com/plugins/deploy/connect-chatgpt>

## Search materials, papers, work and events

Example: “Find recent experiment reports and open GitHub tasks for my project and summarize them with source links.” Specify the project, period and keywords. Search readable titles, bodies, events and tasks, then fetch selected sources. Check observation times and coverage. Missing results do not imply inactivity. Reads do not use LUNA generation allowance or grant access to another person’s private sources.

## Daily planning, editing and PI sharing

Example: “Read today’s notes and propose three goals. Show the original and changes before saving.” Read the current day and revision, then review and save goals, results, blockers and next steps. Saving is not external delivery. PI sharing separately reviews the grant, recipient and exact body. Mentions and document citations provide context, not permission to send messages. Per-item actions retain the saved item and exact revision.

## Research goals, progress, literature and dossiers

Example: “Prepare a paper goal that can be validated in ten days using the selected project. Keep automatic analysis off.” Review and save goals, annual categories, actual submission/acceptance/publication dates, personal automatic caps and manual reserves. Strategy generation uses personal allowance. Read saved proposals and versions; start, pause or resume a detailed dossier. Review linked outputs and automation state when trashing/restoring a goal. Only platform administrators change the shared public-analysis budget. Status checks and downloads do not require regeneration.

## Use research context, papers and figures

Example: “Propose two ways to adapt the architecture on page 6 of the paper attached to my goal, with ablations.” First read the goal’s context, scope, constraints, sources and attachment readiness. Explicit document citations are prioritized; mentions identify people only. A ChatGPT upload is not a LUNA upload. If the current connection lacks a file-transfer tool, attach the file in the platform goal editor and continue using the goal ID. Goal context reads include prepared text passages; `visual=true` also returns the selected page image. Analyze extracted PDF text and the selected visual page without claiming to have read every figure.

## Code work, implementation, ablations, checks and GitHub

Example: “Prepare a new code workspace from this proposal and design ablations that preserve the data loader.” Confirm the repository, base branch, goal and intent. Multiple workspaces may use one repository. Creating one only saves a record; it does not commit. The workspace title later becomes the commit and PR title. Attach files through the platform code editor; `consultant_code_read` returns preparation states and extracted evidence, and `retry_inputs` retries failed extraction without AI usage. Read code → choose a recommendation → implement/ablate → inspect files, documentation and checks → preview publication. Reuse other saved recommendations without regenerating. CPU checks are actual bounded tests, not GPU training or proof of improved performance. Imported local results are labeled owner-provided. A new branch/draft PR is the default; direct base-branch commits require review of the exact diff and baseline SHA. File/ZIP links require platform access.

## GitHub issues, comments, files and PRs

Example: “Draft a comment asking for readiness on this Issue. Do not send it yet.” Select an authorized repository or issue, save a draft, and review the exact destination, body and acting account. After approval, create an issue, comment, file commit or PR. Check IDs and links because titles may repeat. If delivery is uncertain, inspect the existing receipt rather than creating another action. Successful delivery does not prove reading or completion.

## Conversations, direct messages, discussions and reports

Example: “Prepare a discussion comment asking about the experiment conditions below this report.” Select the conversation, direct recipient, material/report discussion or PI share destination. Review the audience and body before approving. Do not copy private source passages to a broader audience. Read delivery and subsequent replies from action history. A mention does not replace destination and preview checks.

## Generate reports, articles and events

Example: “Prepare last week’s report, private to me, without email.” Confirm grant, period, format and audience for reports; date, time, timezone and participants for events; and lab, period and topic for articles. Supported direct tools queue generation after preview approval. Website handoffs must be completed on that page. Article generation is not public publication. A queued link is not a completed body; check readiness and the actual result.

## Report requests, responses, manual work and follow-ups

Example: “Prepare this proposal as manual work and a follow-up tomorrow.” Manual work records a private action without external publication. Report requests and responses may continue in prepared website forms. Enter results yourself or explicitly refresh linked responses; later sharing via a report/message requires a separate reviewed action. Keep self-reports, replies and your own follow-up confirmation distinct. These records inform the next AI review without declaring research success.

## Notices, support and website handoffs

Example: “Prepare a bilingual feature announcement draft and a review link. Do not publish it.” Read permitted workflows and targets, then prepare a website review link. Notices require administrator authority and publication review on the website. Support drafts should include the problem, affected screen and attempted fixes, excluding secrets and research content. A link is not a published notice, submitted support request or delivered report request.

## Costs, progress, recovery and disconnection

Help, saved-result reads and ordinary previews do not invoke new LUNA generation. AI recommendations, drafts, strategy and code use the displayed personal allowance; ChatGPT usage is separate. An AI recommendation “prepare” operation may spend allowance, so review its cost. Distinguish queued, preparing, analyzing, complete and failed, and re-read the same job ID. Never automatically retry an uncertain remote write. Web status refreshes about every eight seconds while active or thirty seconds while idle, pausing in hidden tabs. Disconnecting prevents future access but cannot recall delivered data. Official file contract: <https://developers.openai.com/plugins/reference>

## Continue improving saved results

Choose Continue discussion on a strategy, code result or report. Give ChatGPT the work ID to read its baseline, recent discussion, confirmed decisions and revisions. Example: “Add only two comparison experiments and preserve the rest.” Revisions are saved separately without overwriting publications or GitHub files. Exact-excerpt patches, previous-version selection and alternative branches preserve earlier work. Code may use a saved implementation/ablation as parent; untouched generated files are retained. Prior passing tests do not validate revised code. Only decisions and revisions saved through tools are shared, not your entire ChatGPT history. Use `consultant_thread_list/read/prepare/commit`; saving does not invoke platform AI.

## Shared capabilities in native AI consultation

Start a private research consultation from the bottom conversation dock: AI consultation, then New conversation. Ask AI in the help panel is for usage guidance, not research execution tools. Research AI consultation is processed independently of external collection. In follow-up work, Discuss answers questions; Create revised version explicitly saves a private Markdown version, without publishing or committing. For long documents, revise selected sections. In text-response mode in an unfiltered private consultation, request the same search, read, preparation and execution capabilities in natural language. One question may use up to three personal AI attempts to retrieve context and reason about it. Model-proposed changes become review cards; publication, commits and settings changes wait for approval. Additional AI generation has separate usage. Shared or source-filtered consultations retain their existing evidence boundary; use an unfiltered private consultation for the full tool set. Permission changes, changed baselines and uncertain outcomes require review, not blind retransmission.

## Recover a daily plan draft

Daily plan/review drafts use the administrator call-time policy. Shared daily allowance, daily research cap, per-request call budget and timeout are separate. Details and recovery shows the applied policy and linked attempts. Rejected responses are encrypted for up to 24 hours; revalidation uses no AI when current permissions and evidence still match, and never applies the draft to personal notes automatically. Retry proposes one additional AI call using current policy/evidence. Replaying the same approved review does not duplicate execution. A timeout without a received response does not claim saved output. Old failed responses that were never retained cannot be recovered. In ChatGPT read `daily_drafts` from `consultant_brief`. Use `consultant_prepare` save with `daily_draft {run,operation:revalidate|retry}`, show the returned review, and call `consultant_commit` only for the approved operation.

## Converse directly in the same research work

Continue directly from a new question or existing work in the space. The stable work URL `/workspace/consultant/works/<work_id>/` opens the same work as the web panel on its own page. Copying the link grants no access; reading the page creates no work, conversation or AI request. When reviewing a particular report, code output or earlier version, preserve its result link and selected conversation and version. `consultant_thread_prepare(operation=save, values={desk:{operation:read,value s:{space_id,tab:home/work/results/schedule/trash,work?,query?,offset?}}})` reads currently authorized work, results and downloads. Use the returned `work_id` or `kind:ID` for work. Status accepts `space_id` and optional work for a small progress response. Read/status need no commit and create no work, conversation or AI request. Existing reports, code and goals retain their original IDs and ACL under the common work registry. Distinct analyses are not merged by title. Current conversation revisions and latest saved checkpoints represent the result; older versions stay in detailed history. Sharing or opening a link is not proof that answers or files were preserved. Never bypass blocked content or downloads.

## Distinguish asking, saving notes and preserving external answers

For desk ask or note, supply `space_id`, question, a stable UUID `request_key` and the existing work. Omitting work creates private new work from the first question; reuse the returned `work_id` and thread afterward. Relations can provide multiple origins and references. Pass the exact reviewed review and `preview_token` to `consultant_thread_commit`. Ask requests a platform AI answer and uses allowance; note saves a private record without another model call. If ChatGPT has already produced an answer or files, deliver them through existing sync `begin/chunk/finish` on the returned thread instead of calling ask again. Use the current `sync.revision`, stable `request_key` and exact file manifest; claim saved only after a ready receipt. Undelivered external conversations are not automatically captured, and private conversations or source access are not automatically shared. Verify actual host tool availability and invocation separately.

## Research work trash and restoring previous cycles

Desk `trash_work/restore_work` takes `space_id`, `work`, its current `lifecycle_revision` as `revision`, and `request_key`. `Trash_space/restore_space/reset/restore_reset` uses the current desk revision, `space_id` and `request_key`; `restore_reset` also needs the target batch. Commit only the exact reviewed target and effects. Trash is not permanent deletion; reset archives earlier work as a previous cycle. Original reports, conversations, files, parent materials, members and external sources remain preserved. Execution is blocked while work is trashed or archived. Restoration itself queues no AI; original automation policy remains preserved, so an enabled policy may run at its next scheduled time after restoration. A call already dispatched to an external provider cannot be guaranteed recalled. Work individually trashed earlier stays trashed when restoring a cycle.

## Common uploads and team research

For new input attachments, use desk

`upload_policy/upload_begin/upload_chunk/upload_finish/upload_read` to preserve up to 50 files totaling 100 MB in chunks. Follow returned `chunk_bytes`, actual sizes and SHA256 hashes, and stable request keys. Poll preparation and supply completed attachment IDs to ask or note on the same work. Generated output files use the same 100 MB/50-file original transport. Without a linked space, use `upload_begin` with `inbox=true`, then pass returned file IDs as `content.stored_files` to `capture.save`. With a space, connect those same IDs to the current work. `resolve(fallback=inbox)` chooses the private inbox when no destination is known; invalid explicit destinations and revoked permissions still fail. The existing 8 MB inline checkpoint route remains a compatibility transport for small files, not the overall output-file limit. Downloads stream preserved originals. Hosts follow `host_chunk_bytes`. Preserving input is not a claim that AI has read every original. Review exact content, files and audience through desk `share_publish/share_withdraw` before commit. Read/ask/note with `audience=team` on the same work and deliver outputs to the returned thread. Do not include private source conversations or personal memory as team evidence. Actual external host tool support and byte transfer require separate verification.

## Research team membership and access

Use `consultant_thread_prepare(operation=save, values={space:{operation:team_read or team_search,values:{space_id,query?,offset?}}})` to read members and eligible candidates. `team_add` accepts users, optional roles, access, revision and `request_key`. `team_update` uses user, roles and access; `team_end` uses user; `team_configure` reviews name, description and parent settings under the same revision and `request_key` contract. Show the exact members and access, then pass the returned review and `preview_token` to `consultant_thread_commit`. Every surface uses the same space ID and rechecks current permissions. Adding an existing member cannot overwrite access; authorship never elevates access. Reading or editing roles invokes no AI and sends no invitation. Copying a report link preserves its identity and language without granting access. Verify actual tool availability in the host separately.

## Refine requests and select suggestions

Use `consultant_thread_prepare(operation=save, values={compose:{operation,values}})` for the shared composer. inputs and suggestions only read sources and saved suggestions within `space_id` and the chosen work/thread/result/audience. Selecting a suggestion neither executes AI nor adopts work. After editing returned question and relations, pass `selection_token` in `desk.ask/note compose_selections` to bind the original suggestion to the single actual saved work. Only when explicitly requested, draft/refine reviews and commits question, `request_key`, `space_id` and selected evidence/attachments. It uses one private AI wording call and starts no research work or external action. Poll `compose.read(space_id,run)`, then preserve provenance with `refine_run` on the final send. Host-written wording can be edited and sent without this extra call. Preserve existing ChatGPT results with `capture.save` instead of triggering another AI request. Verify actual host tool availability and delivered file bytes from receipts.

## Hand off the current request and attachments to ChatGPT

Continue in ChatGPT prepares copyable text from the final edited question and selected materials in both new and existing work. Read `compose.handoff` with `consultant_thread_prepare(operation=save, values={compose:{operation:handoff,values:{space_id,request_key,question,work?,thread?,result?,relations?,attachments?,compose_selections?,refine_run?,audience?,language}}})`. `question` is the current input and `attachments` contains verified completed-upload IDs. Never represent filenames or undelivered host files as preserved attachments. Current permissions and the exact work, conversation and output versions are rechecked, and preparation state is returned. Handoff is read-only: `saved=false` and `ai_used=false`, with no work, conversation, answer or external transfer created. Keep `request_key` for retries and choose the destination before creating new work. Do not claim analysis before actually reading the originals. If host tools are unavailable, continue on the website and do not claim preparation, delivery or saving occurred. Use existing `capture/sync` to preserve actually delivered replies/files and verify the completed receipt. Wording refinement, platform AI execution, broader sharing and publication have separate costs and approval boundaries.

## Work maps and exact report evidence

The default graph folds strategy reports into their owning work. `outputs=true` explicitly expands exact report and code-result cards. Renaming changes the work display name without rewriting report content, versions or evidence, or duplicating work. Derivation and reference remain separate relations. When multiple report versions or selectors appear as one work link, its relations retain original source/target keys, hashes, revisions and relation IDs. `scope` describes currently displayed works, outputs and materials within a partial view, not collection totals. Never merge work IDs by list order or title. Read `links/link_history` for detailed relationships and deliberately refresh evidence only after current-access and exact-version checks. Reading the graph makes no AI call.

## Search, inspect and export space literature

Read bibliography through `consultant_thread_prepare(operation=save, values={literature:{operation,values}})`. `list` accepts `space_id,q?,page?,limit(3/20/50)`; `detail` accepts `space_id,id`. Search returned literature:UUID keys in `compose.inputs` and use them as reference relation sources so web, native AI and ChatGPT use the same versions and permissions. Example: “Find contrastive-learning references and pinned Zotero papers in this space, without generating new summaries.” `export` accepts `space_id,format(bib/ris/json),scope(all/search/selected),q?,ids?,exclude_ids?`. Use exact IDs for selected; use `scope=search` with `exclude_ids` to select all matches except specific papers. Use returned content, filename and count unchanged, and provide downloads only when the host actually creates a file. All means the entire authorized range, not the first page. Examples: “Create BibTeX for only the three selected papers” or “Export every search match as RIS and report the count.” This service performs no remote Zotero writes or PDF uploads.

## Edit bibliography and distinguish abstracts from AI summaries

`literature.edit` accepts `space_id,id,revision,changes`. `changes` uses CSL-JSON fields such as `title,author,issued,container-title,DOI,URL` and `abstract`. Use `given/family` for people or `literal` for organizations, and `issued.date-parts` or `literal` for dates. `null` removes a space override and restores the source field. Review the exact before/after and pass returned review and `preview_token` unchanged to `consultant_thread_commit`. Example: “Correct only the publication year, preserving the Zotero original and earlier reports.” Editing names or DOI retains paper IDs and work references. Review `pin(id,revision,active)` under the same permissions. `summary` uses one AI request to store a summary and a separate space-purpose relevance assessment. Read `summary.relevance text/category (direct, partial, unclear)` and `stale`; preserve the original source assessment. It accepts `id,revision,request_key,language(ko/en),regenerate?` and queues a platform AI request only after explicit review and commit. Never call `summary` on `browse, collection, pin` or `export`. Prefer an original abstract; body summaries require accessible preserved text. Original abstracts remain separate from AI summaries and research relevance. Answering in the host differs from requesting LUNA generation. “Show the abstract” is a read; “Request a LUNA AI summary and check completion” is an explicit metered generation.

## Same-work continuation, derivation and kind-specific conditions

Use `continuation_mode=same` or `new` in `desk.ask/note` and `compose.handoff/native_handoff`. `same` appends `request/result` pairs to the same work; `new` creates derived work from the selected originating work/output version on actual save. Preparing or copying creates nothing. Reuse `request_key` for retries and never duplicate work by result or title. Saved requests retain a settled default-condition snapshot. Work-menu conditions and conversational refinement share kind-specific fields while applying a draft and changing defaults remain distinct. Flow sorting places downstream causal work first without turning references into execution dependencies. Literature, code or training classifications and saved conditions never authorize GPU execution or GitHub commits. Distinguish “Check another assumption inside this work” from “Create separate implementation work from this exact output version.”

## Evidence from papers to implementation and future execution limits

Inspect `source_proofs`, current hash and `code_links` for source versions and code provenance. A public code URL proves neither official implementation, reuse permission, environment readiness nor successful execution. Model-design and implementation requests should connect paper IDs, chosen methods, pinned source commits, target repositories/branches and validation plans. Example: “Verify the paper and official-code evidence, then propose only changes for our project; do not commit or train yet.” Always verify which implementation, upload and execution tools the host actually exposes. Preserve only actually delivered results/files through capture; a source URL is not a preserved downloadable file. Training, server execution and independent reproduction have separate capabilities and authorization scopes.

## Connect reference registration to work history

When adding references in a linked work, pass `context={work,thread?,question,channel:chatgpt,audience,continuation:same}` through `consultant_material_prepare` or the existing `consultant_thread_prepare` save/materials envelope. Registration and its delivered request/result are saved atomically; reuse `request_key` for retries. If the destination is unclear, use `capture.resolve` for current bindings and concrete candidates, then ask the user. Never guess by title. Unresolved details may be preserved in a private inbox draft without executing or sharing research. For an earlier unlinked receipt, review materials `operation=link_history` with its `request_key` and context. Do not register papers again or reconstruct unseen conversation. Example: “Link the registration I just completed to the current literature-review work.” To attach a PDF, review literature `operation=attach_pdf` with values=`{space_id,id,revision,request_key,file_id,version_kind}`, then commit. `file_id` must identify a completed upload, not a filename or temporary URL. Use published, preprint or supplement. Example: “Attach this newly found PDF to the existing reference and preserve earlier files.” Native private consultation supplies its linked work and actual question. Team consultation uses fixed shared-work tools and evidence accessible to every current participant. Report only saved text and actually transferred files as complete; distinguish unavailable transfer and partial completion.

## Complete server tool inventory

For diagnostics. Host discovery must be verified separately.

`search`

`fetch`

`source_connection_status`

`search_work_items`

`list_my_generation_jobs`

`list_workflows`

`prepare_workflow`

search\_materials

search\_help

read\_help

action\_options

prepare\_action

get\_action\_status

quick\_help

render\_quick\_help

create\_github\_issue

create\_calendar\_event

generate\_report

generate\_article

consultant\_brief

consultant\_prepare

consultant\_commit

consultant\_strategy\_list

consultant\_strategy\_read

consultant\_strategy\_prepare

consultant\_strategy\_commit

consultant\_code\_list

consultant\_code\_read

consultant\_code\_prepare

consultant\_code\_commit

consultant\_thread\_list

consultant\_thread\_read

consultant\_thread\_prepare  
consultant\_thread\_commit  
consultant\_account\_read  
consultant\_account\_prepare  
consultant\_account\_commit  
consultant\_account\_usage  
consultant\_account\_contributions  
consultant\_account\_diagnose  
consultant\_material\_prepare  
consultant\_material\_status  
consultant\_material\_commit  
consultant\_work\_read  
consultant\_work\_status  
consultant\_literature\_list  
consultant\_literature\_read  
consultant\_literature\_evidence  
consultant\_literature\_export  
consultant\_literature\_edit\_preview  
consultant\_literature\_metadata\_preview  
consultant\_literature\_visibility\_preview  
consultant\_literature\_pdf\_preview  
consultant\_literature\_summary\_preview  
consultant\_literature\_commit  
consultant\_files\_import  
consultant\_files\_status

consultant\_files\_import\_pdf\_url  
consultant\_capture\_resolve  
consultant\_capture\_save\_preview  
consultant\_capture\_commit  
consultant\_capture\_read  
consultant\_zotero\_destinations  
consultant\_zotero\_prepare  
consultant\_zotero\_commit  
consultant\_zotero\_status  
consultant\_workspace\_capabilities  
consultant\_workspace\_status  
consultant\_workspace\_prepare  
consultant\_workspace\_commit  
consultant\_journal\_options  
consultant\_journal\_prepare  
consultant\_journal\_commit  
consultant\_journal\_status  
research\_action\_options  
research\_prepare\_action  
research\_action\_list  
research\_save\_draft  
research\_preview\_action  
research\_action\_status  
research\_refresh\_results  
research\_update\_followup

research\_post\_github\_comment

research\_create\_github\_issue

research\_send\_chat

research\_send\_direct

research\_post\_discussion

research\_share\_pi\_report

research\_commit\_file

research\_create\_pull\_request

research\_record\_manual\_work

research\_open\_report\_form

research\_open\_calendar\_form

research\_open\_request\_form

research\_open\_response\_form

research\_open\_code\_workspace

research\_open\_report\_editor

teaching\_options

teaching\_prepare

teaching\_commit

teaching\_status

teaching\_list

teaching\_cancel

Use only tools available in your current connection.

## Continue saved analysis and intermediate results

Examples: “Show the original request and saved parts”, “Resume the failed stage.” Use `consultant_strategy_read` with `status/history/analysis`; `prepare` and `approve analyze` or `analysis_pause/analysis_resume/analysis_cancel` before changes. New analysis defaults to four stages with a four-attempt cap. Web and native consultation share the same IDs and outputs. Reads use no platform AI; generating a conversational answer may use allowance.

## Personal settings through ChatGPT and AI consultation

Ask “Show my remaining AI allowance and reset time”, “Show my contribution score”, or “Show my administrator-managed execution policy”. Shared account tools read current settings and prepare an exact before/change review. Execution policy changes require current platform administrator authority and explicitly review the effect on all users. Personal overrides are rejected. Other settings remain personal; sharing rights and publication remain separate. Tools: `consultant_account_read`, `prepare`, `commit`, `usage`, `contributions` and `diagnose`. Help AI can read current personal settings and usage. Continue to work/settings consultation carries the question into private full consultation; subsequent reasoning uses personal AI allowance and changes require review. If the ChatGPT host does not expose a tool, check the connection and granted scopes.

## Analysis continuity and ChatGPT

New strategy analyses default to prior reports, feedback, saved follow-up decisions and recent discussions for the same goal. Originals remain available; bounded excerpts and omissions are recorded for model input. Independent analysis excludes prior analysis context. Revoked evidence is never reused. Code consultation from a report, section or action retains source provenance. ChatGPT resumes the same request ID and checkpoints; completed work uses follow-up analysis. `analysis` with `status=true` returns small metadata. Review `strategy_prepare` acknowledge with `kind/id/version` before `commit`. If the host exposes an old limit or missing operations, refresh the connected tools and read the central policy.

## Synchronizing external discussion and code results

Each code workspace has one representative notice across preparation, AI and validation steps. Dismiss preserves failures and does not mark them successful. Link a strategy report under Sources and links; the report also links back to derived code. External results preserve the question, answer, files, provenance, versions and downloads. Ask ChatGPT or Codex to continue this LUNA workspace and save the answer and files to the same workspace. Only a ready storage receipt confirms synchronization. Conversations or temporary files never delivered to LUNA are not automatically imported. You can also paste an answer and select files in the platform. Conversation text is limited to 120 KB, files to 100 files and 8 MB total; retry with identical files and request identity to resume. Files are not executed; model and test claims are producer-reported. Revisions preserve previous results. Storage uses no new AI call but consumes storage and backup capacity. Encrypted rejected responses may be revalidated without AI for up to 24 hours; invalid responses remain failures. Execution/resume settings start collapsed.

## Calendar ranges and official reference prices

Today/this week/this month/this year set date bounds in Asia/Seoul. Choose chart granularity separately in advanced filters. Administrators receive OpenAI Standard reference defaults checked on 2026-10-05. Custom rates and historical estimates are preserved. Long context, cache writes and processing tiers are distinguished; unknown tiers use an API-equivalent estimate, not subscription billing.

## Active work, archives and deletion

The live strip shows queued/running requests only. Analysis archives separate attention, completed, cancelled/failed and trash; saved outputs remain downloadable. Cancel before moving active work to trash; restore never resumes automatically. Permanent deletion requires an exact scope review and current permission, and rejects references from later consultations/code. Original attachments and usage totals remain. Old encrypted backups are not immediately erased; reapply the latest deletion ledger after restoration. ChatGPT uses `consultant_strategy_prepare analysis_trash/analysis_restore/analysis_purge` and `commit`.

## Complete literature and report history

Goal settings precede research context. Paginate all literature/search observations and reports. Deduplicated literature and per-search snapshots are separate; legacy observations have unknown original search batches. Only the final analysis stage creates a final report; partial results remain on the request. Compare two saved reports without AI. ChatGPT `strategy_read` collection=`reports/literature/observations` with `goal,page` reads the same lists.

## Start repository consultation

Choose a project, repository and purpose using Start GitHub repository consultation on the code landing page. Continue an existing analysis where available. Consultation is expanded on the detail page and requires a completed repository read. Connection/permission problems link to connection settings. The current bound of 24 text files/85KB does not establish whole-repository verification.

## Evidence access and analysis recovery

Analysis detail shows unavailable evidence counts and Google reauthorization status. Open the linked connection, personally authorize Google, wait for collection, then check access again. Authorization alone never restores historical access. Reuse saved stages only after a new authorized collection confirms the same file, version and content. Changed content or revoked access keeps old results preserved but restricted. Review a new analysis with accessible evidence to create a separate four-stage request after reviewing the original question, budget and timeout. Blocked source text and old outputs are not copied; new calls consume the existing personal allowance. Preview/status reads never run AI. ChatGPT uses `strategy_read` analysis diagnosis and the exact `strategy_prepare analysis_recover / strategy_commit` review contract.

## Research daily limits and the overall allowance

Use Manage limits in Today's research to open Daily usage limits in Platform AI management.

Administrators choose Follow overall AI allowance or Set a separate cap. Administrators inherit by default; the existing regular-user default remains 5. Inheritance is not extra credit: research shares the personal pool with translation, help and other AI. A separate cap of 0 disables research AI.

Manage my overall daily AI allowance opens per-user overall limits. Attempted runs, queued reservations, completed and failed runs and reset times are distinct.

The overall pool resets at Seoul midnight; a separate research cap resets at midnight in the research timezone. Per-request budgets and per-call timeouts remain separate execution settings. Raising limits never clears history or usage, or automatically retries failures.

In ChatGPT or AI consultation, read `research_usage` via `consultant_account_read`. Only administrators can review and commit execution changes to `research_daily_limit/admin_research_daily_limit`; null inherits the overall pool.

## Enable or disable automation for a goal

Use the automation control at the top of a paper goal without deleting it. Conditions are collapsed by default. Periodic review offers None, daily, every three days, weekly and fortnightly. None disables periodic review only; new literature, changed evidence and deadlines within seven days are independent triggers. Disable automation to stop every automatic trigger. Literature discovery is separate and does not consume personal AI attempts. Disabling cancels unstarted automatic requests; a started call may incur its existing cost and save a validated result before stopping. Manual analyses, reports, literature and history remain. Enabling or changing conditions schedules future work without replaying missed periods. Allowance waits, failure holds and running states are distinguished; the next review time is an estimate. Ask ChatGPT or AI consultation to disable automation for the goal. `consultant_strategy_read` returns automation and its revision; `prepare(operation=goal, values={automation_revision, automation:{automatic:false}})`, review the exact change, then commit updates the same policy. Actual host tool exposure must be verified separately. Reading and saving policy do not execute an AI analysis.

## Preserving accumulated literature and search history

Editing goal settings preserves discovered papers, analysis links and observations. Titles and deadlines do not change the search revision; changing search terms does. Accumulated literature is the default, with filters for current, earlier or unknown settings, recent discovery, analysis availability and relevance review, paginated by 20. Unknown historical search terms are never inferred. Prioritize, select by relevance and exclude apply to future analyses without deleting previous results. Reassess relevance after changing the question or research scope. Reads and selection saves spend no AI allowance; automatic analysis consent is separate. `consultant_strategy_read(goal)` returns the same library in ChatGPT and AI consultation. `prepare(operation=goal, values={literature:{operation:list,filter:all,page:1}})` is read-only. For changes, use `literature operation=select, paper, inclusion(auto/include/exclude)` and revision, review, then commit. Search observations retain the actual terms and settings.

### 3. Sharing, Notion and printing

Use the top share icon for team sharing or Notion publication. Find recipients by name and use ? for field help. For Notion, select an allowed parent and review the body before queuing publication. Configure a connection if none is available. Notion audiences are independent. The export icon provides full-body print, PDF and Markdown. An uncertain publication is reconciled, never blindly reposted.

### 4. ChatGPT and AI consultation

Connected tools support reading a goal workflow and recording notes or imported results. Use the lifecycle envelope inside consultant\_thread\_prepare save, followed by reviewed commit for changes. Only transmitted content is saved. Reads and record writes do not invoke extra AI or remote experiments; new consultation follows existing usage policy.

#### 4. Continue across AI surfaces

1. Use beside the send button in either a new or follow-up request to choose: Answer in this work, Continue in AI chat, or Continue in ChatGPT.
2. Selecting a destination runs no AI. Answer in this work sends the request. Continue in AI chat prepares your final question, selected evidence and files in the consultation panel; review and send there to start.
3. Continue in ChatGPT prepares handoff text. Copy it into ChatGPT with LUNA connected. Copying alone neither transfers nor saves anything externally.
4. Reuse the same work and conversation for the same purpose. New work is created on actual send; switching a prepared request's destination must not run it twice.
5. Only delivered and saved content and files are preserved. Check the receipt's work ID, files and download links. Choosing a destination does not broaden private conversation access or sharing.

### 7. Core materials and AI drafts

1. Search and pin important sources in Materials & results. Pinning and unpinning do not copy, delete or change source access.
2. Private pins enter only your own AI context; team sources still require current source permissions. Attach relevant evidence and choose Ask AI to save a request on the same work.
3. Choose Save note to preserve a private memo without AI; it does not replace the current result. Waiting, running and failure reflect actual request state; saved stages remain preserved.
4. AI requests follow administrator policy and personal limits. Reading, saving notes and pinning make no new generation call.
5. Pins and AI drafts are not factual verification or researcher approval.

## 10. Find earlier analyses and results

1. Existing goals, code, reports and discussions retain their original IDs and permissions while being linked into one work list. Open work or search Materials & results for originals and downloads.
2. The current revision and latest saved checkpoint represent a conversation; earlier versions remain in detailed history. Distinct reports are not merged by matching titles.
3. If evidence is inaccessible, the owner may see only a preserved-result notice. Restore current connection and permissions before opening content or downloads.
4. Never infer or claim to have read blocked content. A classification-only change with identical content, version, connection and current access does not require regenerating saved results.

## 11. One work list and research map

1. Choose List or Research map in Research work. Selecting work opens the right panel. Use ↗ or Ctrl/⌘-click to open the same work on its own page with a stable URL. The default map groups results inside their owning work; Expand outputs reveals exact report and code-result cards.
2. Renaming work changes its display name without rewriting source titles, content, versions or evidence, or duplicating work. Distinct works are never merged by title.
3. Solid workflow/derivation links and dashed reference/validation links are both visible by default. Multiple origins, references and sections retain their exact report IDs and versions even when output cards are folded.
4. Counts distinguish the works, outputs and materials currently shown, not the total saved collection. Use the next work range or a focused neighborhood to navigate.
5. Renaming or reordering the list does not reposition the same works and relationships. Goal settings, literature, strategy analysis and specialized code tools remain available from that work's settings or detail screen.
6. Source changes require deliberate evidence review and refresh. Disconnecting does not delete the source or earlier history.
7. Maps, search and opening a work URL call no AI and create no work or conversation. The panel and separate page use the same work, current permissions, conversation, results, files and history.
8. While drafting, use ↗ to open a new tab. Unsent text and selected files remain in the original tab. Open in research space returns to that work. Use the original tab to resume its search, map position and draft.

## 12. Trash and a new research cycle

1. Review the target and effects before moving work or a space to Trash. Saved conversations, outputs, attachments and external originals remain preserved.

2. A new cycle archives earlier work while retaining materials and members. Restore from Trash or previous cycles without automatically restoring work individually trashed earlier.
3. Related local pending or running requests stop, but a call already sent to an external provider cannot be guaranteed recalled. Restoration itself queues no AI.
4. Original automation policies remain preserved, so an enabled policy may execute at its next scheduled time after restoration.

### 13. Navigate quickly and inspect history

1. Switch between List and Map and sort by recent activity, creation, execution state, attention or title. Changing work loads its detail in place and retains the draft question for each work.
2. Open Records and filter All, Conversation, Results and versions, Executions or Links. An answer and its result are grouped.
3. A conversation started from an earlier report or selected passage retains its original thread. Load earlier records on demand.
4. No execution requested, generating and needs attention describe execution separately from saved outputs. Downloads require actually preserved files.

### 14. Save naturally from a conversation

1. Ask the AI consultation or ChatGPT: "Save this conversation and its results to the current research work." LUNA prioritizes an explicit work, the bound conversation and a previous save receipt. Choose by title when several targets match, or preserve privately in the personal inbox before selecting a destination.
2. New work requires an explicit request; retrying the same request does not duplicate it. Review the destination, content and files before saving.
3. The personal inbox supports downloads and linking to existing work. Saving text makes no additional platform model call.
4. Files can be uploaded in resumable chunks up to 8 MB per bundle; larger materials may be saved as accessible source references. A source link is not a preserved binary copy.
5. Check the completed save receipt: uploading means only received parts are preserved. External messages and files that were never delivered are not automatically collected.
6. Team sharing and external publication retain separate scope reviews.

### 15. Full-width map and images

1. The map uses the full work area. Derivation and dependency relations guide its left-to-right layout, including independent work and loops.

2. Select a node for detail and return to the same map position. More offers PNG, JPEG and SVG exports of the currently loaded range, not a claim to cover every record.
3. No AI call is made.

## 16. Attach files and use code results

1. Use Add materials for files, images and platform evidence; paste links into your question. Upload up to 50 files totaling 100 MB through resumable chunks.
2. Originals are preserved while AI excerpts are prepared in the background. Preparation is not a full read of every original; extraction failures and text or image limits are shown separately.
3. Actually completed ChatGPT file transfers can be attached with the same permission checks. Missing code output has no download; execution history shows the failure.
4. Import preserved text code or a validated ZIP as a code changeset and inspect each diff. ZIPs are never executed or automatically installed.
5. Review repository, branch, baseline commit and changes before approving GitHub publication. Reconcile uncertain responses to avoid duplicate posts.

## 17. Keep existing work in My saved items

1. Star existing work, conversations or results to keep a shortcut in My saved items. This differs from a preserved copy of delivered conversation content and files.
2. Shortcuts reopen the same work and version; source permission revocation blocks content access. Web, AI consultation and ChatGPT use the same bookmarks and save receipts.

## 18. One library and pinned materials

1. Use All, Favorites and Saved conversations and files in one personal library. Old favorite URLs and folders remain valid.
2. Unstarring does not delete preserved originals. Materials and results initially shows pinned items; search includes all accessible materials.
3. Toggle the same pin to unpin and clear search to return to pins. Pinning does not expand source permissions.

## 19. Private work and team discussion

1. Click Only me on a work item to manage sharing. Review the exact body and selected files before sharing.
2. Earlier private discussions and personal memory are not shared automatically. Continue in team view to accumulate attributed discussion under the same canonical research work.

3. Sharing withdrawal, membership and source access are rechecked; private originals remain preserved. ChatGPT uses the same work ID and current permissions and reflects only delivered content.

## 20. Add schedules and export maps

1. Use Add to choose an event, GitHub task or report request and then create or link existing. Creation opens the standard form with optional AI drafting.

2. Linking does not duplicate the original or create an external item. Select one image format before downloading a map.

3. Dashed reference and validation relations identify evidence or verification links, separately from derivation order.

## 24. Refine a request through conversation

1. Write a few sentences in New research or Continue this work, then select sources and files in Add materials. Put links in the text. Both composers use the same refinement, attachment and destination controls.

2. Refine request opens a conversation on the right. It checks the current request, work and selected evidence, suggests a research kind and identifies missing conditions. It reuses known conditions and asks one or two important questions at a time.

3. A choice fills the reply box. Add or edit details, then choose Send reply. You may leave a condition undecided or ask for a recommendation. Repository candidates are authorized connection metadata, not proof of write permission or server availability.

4. Preview the prepared request and conditions, distinguishing your input, existing evidence, AI suggestions and unknowns. The composer changes only when you choose Apply to request. Undo restores your earlier input.

5. Applying does not execute research. Edit the result if needed, choose this work, AI consultation or ChatGPT from the destination menu, then send. You may also skip refinement and send directly.

6. Closing and reopening preserves the saved conversation. On mobile, closing returns to the original composer. An older response never automatically overwrites text you changed.

7. Starting refinement and sending each additional reply request one personal AI attempt each. Page reads, saved conversation reads and applying use no AI. Closing the panel does not cancel an already accepted AI job.

8. Administrator policy, remaining allowance, provider and time limits apply. Failures retain input and conversation without automatic paid retries. Up to 50 files totaling 100 MB are supported; uploading is distinct from analyzing all original content.

9. Selecting a saved suggestion only fills the request and evidence; Save memo stores a record without AI. Team requests use only currently authorized shared evidence.

## 25. Continue request drafts and results in ChatGPT

1. You can hand off a rough request without finishing refinement on the website. Choose Continue in ChatGPT from the destination menu for either a new request or a follow-up.

2. The handoff includes the latest composer text, selected sources, prepared attachments and existing work, conversation and result versions. An applied refinement also carries the same request session and accepted revision. Preparing and copying the handoff calls no AI, creates no work and transfers nothing externally.

3. Paste the text into ChatGPT with LUNA connected and check that the host exposes the needed tools. Ask: “First clarify the request by asking for missing conditions; proceed after we settle it.”

4. Preserve the same request session, work and evidence versions when saving a request refined in ChatGPT. Saving externally prepared requests and results does not rerun platform AI. Draft storage, applying a request and executing research remain separate actions.

5. Ask: “Save the discussion delivered so far and the generated files to this work, then show download links.” The platform preserves only conversation content and files actually delivered through its tools, not the entire external chat automatically.

6. A filename or temporary link is not proof of preservation. Check the receipt destination, file list, downloads and readiness. Use web upload if the host cannot deliver files.

7. Confirm a destination for new work and retain the existing work, conversation and exact result version for follow-ups. Continue through the stable work link; include the exact result or version link when reviewing a specific output. Opening a link does not save drafts, external conversations or files. Private conversations and evidence do not become team-shared automatically. Publication and paid GPU execution have separate scopes.

## 33. Automatic recommendations and usage

1. Open space settings in Research team to turn automatic research recommendations on or off. Existing and new spaces default to on, preparing suggestions from existing reports and consultations.

2. Opening pages, reading suggestions, searching and sorting make no new AI call. Relevant saved results refresh recommendations in the background; the first view shows saved suggestions.

3. Additional automatic AI enrichment defaults to zero attempts. Set a daily cap if needed; the space owner’s allowance and administrator policy also apply. Changing settings never resets consumed usage.

4. Turning off stops subsequent automatic recommendation generation while preserving existing suggestions and history. Pending work rechecks the setting and permissions immediately before execution.

5. Private consultation or source recommendations are not automatically shared with other members. Changed or revoked evidence requires revalidation. Preparing a suggestion does not execute research, prove success or constitute human review.

### 34. Compact cards and work search

1. Open titles in the Projects and Research spaces sections of a grant or project. The directory uses the same cards without repeating a parent already shown.
2. Work and result counts use currently accessible entries. A work's conversation, report and execution history are not counted as separate work. A + marks a bounded count.
3. Choose List or Map and type to search Research work. Search waits for composition to finish; Enter searches immediately. Use the compact sort menu for recent activity, creation, status and other criteria.
4. Searching and changing views preserve the draft and selected work. Renaming work or adjusting map relations retains the same work ID and history.

### 35. Research kinds, requests and results

1. A research kind describes the activity: planning, literature review, model design, experiment design, code implementation, training, validation, reproducibility, analysis, discussion, writing, reporting or scheduling. Reports, code files and literature tables are result formats.
2. AI proposes a kind during refinement. Direct requests receive an initial classification; use Change research kind in the work menu to correct it. Manual choices are never overwritten automatically.
3. Every work uses Request, Results and Continue this work, in that order. Expand long requests and open or download results. Find initial requests, conversation, execution and earlier versions in History under the work menu.
4. Changing kind preserves sources, results and permissions; it does not authorize server execution or GitHub writes. Distinguish a training plan from actual training and a validation plan from completed validation.

### 36. Continue earlier goals and literature

1. Earlier paper-goal settings and discovered literature can be organized as a separate preparation work. Read its research objective and conditions as the request, and open the literature bundle as its result.
2. The literature result includes search and observation history, selection state and sources. Open or download it, distinguishing reorganized existing records from new AI analysis.
3. Earlier strategy and detailed-report IDs, conversations, reports, citations and downloads remain intact. Linking preparation as the logical starting point never invents historical runs or conversations.
4. Earlier settings and literature URLs lead to the preparation work. Its automatic execution settings retain the earlier policy; reorganization does not resume stopped AI or incur additional AI charges.

5. Renaming work never creates another map node for the same work. Work and result artifacts remain distinct; expand result versions when needed.

## Recommend, select and pin core sources

AI source suggestions in Materials & results works without pins. The website reads saved suggestions first; only generating fresh ones requests one AI attempt. In ChatGPT, first verify that this host exposes `consultant_thread_prepare`, then read with `operation=save`, `values={material_recommendations:{operation:read,values:{space_id,run?}}}`. Read returns currently authorized candidates without AI. Inspect title, reason, coverage, provider, url and pin\_state. Prepare `operation:recommend` with `space_id`, `request_key` and optional language; review source scope and the expected one AI attempt, then pass the returned `preview_token` and review to `consultant_thread_commit`. `Operation:study` with a sources list of material IDs prepares a composer link with evidence and invokes no AI. Pinning uses the existing `space.pin` contract for the current private/team scope. Selecting, pinning and sending research are separate actions. Never claim support for unavailable host tools or that unread full text has been checked.

## Start with a short question

First question

Explain the attached paper's main question, method and limitations in plain language.

Select the paper in Add materials before sending.

Clarify an idea

I want to study how to reduce bias in recommender systems. Suggest three questions to narrow the idea.

Leave unknown data and equipment as questions.

Refine wording only

Refine "I want to improve a model with little data" into a request with a clear objective and expected output.

Refine through conversation, then Apply to request. Starting and each additional reply use one AI attempt; research execution is separate.

Save a note

Save this as a private note in the current work: first check uncertainty caused by the small sample.

Saving a note makes no additional platform AI call.

One next step

Identify the most important uncertainty in this work and recommend one next step with a reason.

Adopting a suggestion and executing it are separate.

Clarify a paper goal together

I want to prepare a paper. First check this space's objective and selected literature. Reuse known research questions, venue, deadline, time and resources; ask only one or two missing questions at a time. Use my answers to prepare a final request covering the literature scope and desired outputs. Do not start research or execution before I settle it.

Use Apply to request to put the accepted draft in the composer. Refinement and actual research execution remain separate.

Prepare a repository-aware implementation request

I want to apply this method to the project code. Show authorized connected repository candidates and help choose a repository and base branch. Ask only for necessary environment, implementation scope, small-example checks and completion criteria. Distinguish actually reading a repository from checking connection metadata. Do not commit yet.

A repository connection does not prove write permission or server availability. Review actual commits separately.

## Find literature and core sources

Find materials

Find materials about contrastive learning in this space and show their titles and sources.

Searching saved materials is a read within current permissions, with no new AI generation.

Recommend core sources

Recommend core sources for this research. Show saved suggestions first, and tell me the AI usage before generating new ones.

Reading saved source suggestions uses no AI; generating suggestions uses one AI attempt.

Pin selected sources

Pin the first and third materials I just reviewed as private core sources in this space.

Check the selected sources and audience; source permissions stay in effect.

Compare papers

Compare the questions, data, methods, metrics and limitations of the three selected papers. Distinguish full text you read from summaries.

A preserved file is not proof that its entire text was read.

Start from selected sources

Use the selected sources to identify unresolved questions and propose the most feasible research direction.

Use Research with selected materials, review the evidence and objective, then send.

Find references across sources

Search this space's literature collection for contrastive learning papers. Include collected and pinned Zotero references, showing each paper once. Do not start a new search run or AI summary.

Read the saved collection under current access permissions.

Read the original abstract

Show the selected paper's original abstract, DOI, authors and publication date. Keep AI summaries separate from the original abstract.

If no abstract is available, say so rather than inventing content from the title.

Explicitly request a saved summary

Request a saved LUNA AI summary and separate relevance to this space's purpose and questions, using this paper's accessible abstract or preserved body. First check source coverage and AI usage, then show the summary when it actually finishes.

A host answer and a paid platform summary request are separate. Queued does not mean completed.

Correct bibliography

I want to correct the selected reference's author name and year. Show the current bibliography first. Apply only the values I supply to this space without modifying the Zotero original.

Review field-level changes and the current revision. Recollection preserves corrections.

Export selected references as BibTeX

Export only the three selected papers as BibTeX, including authors, year, DOI and original abstracts. Do not substitute AI summaries for abstracts. Tell me if you cannot create a real downloadable file.

Claim a download only when the host actually creates a file.

Export all search matches

Export all "graph learning" matches in this space's literature as RIS across every page and report the count. Do not upload anything to Zotero.

Export is a standard bibliographic snapshot, separate from Zotero synchronization.

Restore the source bibliography

Restore only the selected paper's publication date to the source value. Preserve other edits and its work references.

A null field removes that override and restores the source value.

Use literature as implementation evidence

Link the two selected papers as references for a new model-design task. Check readable source coverage first, distinguishing applicable method details from unknowns.

Retain paper IDs, versions and permissions, and state what source text was read.

Register references as materials

Register these DOIs and paper links as references in this research space. Reuse existing materials and check the destination project/grant and audience. Do not generate new AI summaries.

Bibliography registration, original storage and AI analysis are separate; registration requires actual host tools.

Register only missing papers

Check each paper from my earlier registration request and register only items not yet accepted. Do not recreate queued materials. Distinguish storage failures from missing PDFs.

Reuse existing request receipts and material IDs.

Attach a PDF to an existing reference

Attach this published PDF to the already registered paper without creating another reference; preserve any earlier preprint. Tell me when it is stored and downloadable, and report text extraction and AI analysis separately.

Verify completion from delivered bytes and the platform preservation receipt.

Find bibliography in Materials

Find the references collected in my research space through Materials. Search by author and DOI and show each paper once after checking duplicate source identities.

Search only currently accessible references without exposing another space's private corrections.

Export corrected bibliography

Check that these bibliography corrections made through Materials also appear in the research space, then export only the selected references as BibTeX. Preserve Zotero originals and other spaces' corrections.

Read and export the same reference IDs and space corrections under current permissions.

Check original and extraction status

Check original-storage status for the registered papers. Distinguish PDFs, archived HTML, extraction failures and missing originals, and show download links for preserved files. Do not start AI analysis.

An archived webpage is not presented as an acquired PDF.

## Work through research stages

Form a hypothesis

Propose two testable hypotheses grounded in the selected literature and observations that could refute each.

Separate source claims from newly proposed interpretations.

Plan an experiment

Design an experiment for the first hypothesis, including data splits, baselines, metrics, repetitions and stop criteria.

Planning does not start an experiment or a GPU run.

Plan implementation

Inspect the linked repository's current structure, then propose small implementation steps and the necessary checks.

Label proposed code, synthetic checks and actual execution results separately.

Interpret results

Calculate changes from the baseline and uncertainty in the attached results, separating supported conclusions from claims needing more experiments.

Select the actual result file and check the calculation evidence.

Draft a manuscript section

Draft a results section using only reviewed findings in this work. Do not invent numbers or citations.

The draft does not automatically change the manuscript source or author approval.

Plan training with unknown resources

I want to prepare training for this model. First check access to the selected data, splits, baselines and metrics. I do not yet know whether a GPU server is connected and available. Do not assume missing resources exist; ask about possible environments and time/cost limits. Save only the execution plan for now; do not start training.

Unverified environments and permissions remain unknown. Saving a plan does not execute a paid GPU run.

Inspect paper-code provenance

Inspect the paper's linked code repository and report whether it is official, the inspected commit, license and environment requirements. Do not commit to my project or run it on a server yet.

Distinguish official links, inspected code, reuse rights and validation results.

Edit saved research defaults

Show this work's research conditions. Change the default evaluation metrics to the ones I confirm while preserving earlier request/result settings. Save settings only, without execution.

Check the kind-specific fields and current settings revision.

## Continue existing work

Follow up in the same work

Analyze the sample-size limitation in this work's latest result in more detail.

Open the existing work before sending to continue its conversation and results.

Continue an exact report

Use section 3 of the selected report version to prepare the next analysis plan.

Keep the selected version as evidence rather than silently substituting the latest result.

Review all suggestions

Show all saved suggestions for continuing this work so I can compare their reasons and expected outputs.

View all reads saved suggestions without a new AI call.

Derive separate work

I want to start separate research comparing another dataset while referring to this result. Preserve the link to the original work.

Keep same-purpose questions in existing work; derive work for a different objective.

Check status and partial results

Show this request's actual state, saved partial results and files available for download. Do not rerun it.

When failed or quota-waiting, inspect saved results and resume conditions first.

Append a request/result pair

Continue within the current work without creating another item. Check two assumptions in the latest result and preserve this request and answer in the same work.

Use `continuation_mode=same` and never attach an earlier answer to a pending request.

Derive from an exact result

Create separate implementation work from the selected output version. Keep the original conversation intact and link this literature and result as evidence.

`continuation_mode=new` creates the new work and originating evidence on actual save/send.

## ChatGPT, files and preservation

Hand off the final request

Check the final question and selected evidence in the copied LUNA request, then continue this work. Check attachment readiness first.

Verify the LUNA tools actually available in ChatGPT. Copying alone does not transfer anything.

Save discussion and result files

Save the discussion delivered so far and the actual generated files to the linked LUNA work, then show download links.

Verify actual content/file delivery and a completed receipt. Saving does not require another AI run.

No work selected yet

Save this conversation and its generated files in my private inbox. I will link them to research work later.

Actual uploads support up to 50 files totaling 100 MB. Filenames or temporary links alone do not establish preservation.

Check original downloads

List the originals preserved in this work with their readiness states and provide downloads for completed files.

Use the website upload when the host cannot deliver the file.

Refine and save from ChatGPT

Help refine the supplied LUNA request in this conversation. Preserve its request session, work and evidence versions, and ask about missing conditions. After saving the settled request, preserve the results I ask you to prepare here and the files actually generated in the same work. Do not regenerate them with platform AI. Show the saved work link, file list and download availability.

The host must expose the required LUNA tools. Check draft/result receipts and actual delivered files; do not describe this as automatic synchronization of the entire external chat.

## Team, relationships and schedules

Rename work

Rename the current work to “Small-data comparison experiment” while preserving its results and references.

Renaming does not create a new work item.

Link several references

Link the two selected works and the report as references for this work, then show their relationships on the map.

Reference, derivation and prerequisite are distinct relationships.

Share selected results

I want to share only this summary and the selected result files with the research team. Show the recipients, body and files first.

Share only the reviewed scope; private conversation history and memory are not automatically exposed.

Link an existing schedule

Link the existing next research meeting to this space without creating a duplicate event.

Reuse the existing calendar or GitHub task; review any external creation separately.

## Find and edit literature in Materials

1. Open References in Materials or search by title, author or DOI. References saved in a research space are discoverable in the same catalog; do not register another copy.

2. Select a card title to open its wide permanent detail page; View all opens the dedicated literature list. Translation, sharing, favorites and print are at the top, followed by authors, venue, publication date, abstract and source links. Original abstracts and AI summaries are separate. Generate an AI summary only with the explicit button or request.

3. Edit only the required bibliography fields. Materials and the space use the same reference ID and correction. Check the applicable research space; other spaces and Zotero originals are not silently overwritten.
4. Select references and choose BibTeX, RIS or CSL JSON under Export. Distinguish selected references, all search matches and the complete current scope. All includes every page in that scope, not just the visible rows.
5. Use File → Import in Zotero to read the exported file. This is a bibliographic snapshot, not automatic two-way synchronization. Unpinning does not delete the original material.

## Register references and request source storage in conversation

1. Tell native AI or ChatGPT the research space and the DOI, source URL or bibliography to register. Example: “Register these papers in this space, reuse existing references and do not generate new AI summaries.”
2. Check the destination project/grant, audience, new/existing references and original attachments. The host must actually expose the required tools. Search-only tools cannot prove registration.
3. Check each returned material link and status. With a linked research work, the actual registration request and receipt also appear in its history. Ask for a destination if work context is unclear; never invent earlier conversation that was not delivered. Bibliography registration, pinning, source storage, text extraction and AI analysis are distinct. Reuse the same request ID when retrying to prevent duplicates.
4. Add a PDF later from the reference detail page. Active authorized participants can contribute files, preserving contributor and file versions. Identical files are not linked twice and earlier PDFs are not overwritten. Distinguish published papers, preprints and supplements. In conversation, identify the existing reference before attaching a PDF. Example: “Attach this PDF to the existing reference without creating another paper.” Only delivered files or accessible originals can be stored.
5. A ZIP or download link generated in ChatGPT is not proof of platform storage. Check file delivery and the platform receipt. Registration, reading, pinning and exporting do not themselves generate AI summaries.
6. The tool sequence is `consultant_material_prepare` → `consultant_material_commit` → `consultant_material_status`. Review the destination and audience, then pass the unchanged `preview_token` and review. A cached host can use the existing `consultant_thread_prepare` save / materials envelope for the same service. If neither route is exposed, refresh the connection tool list.
7. A request accepts up to 50 items and 100 MB of combined attachments. Distinguish “register bibliography only” from “also store an accessible PDF.” A paper landing page is not reported as a saved PDF; supply an actual PDF file when needed.

## Understand storage, PDF and extraction states

1. Once registered, bibliography can be searched and edited before source collection finishes. Do not create duplicate references because earlier files are still processing.
2. Queued and running storage are distinct. Stored means a verified preservation receipt exists. An archived HTML page is not a PDF.
3. If the original is stored but text extraction fails, it remains downloadable. Combined source attachments may total 100 MB; safe text extraction supports up to 20 MB per file. Larger originals are preserved with an extraction-limit status. Scans, protected documents or unsupported formats may not yield readable text.
4. Retryable storage failures use bounded attempts and backoff. For persistent failures or access requirements, inspect the reason and supply an accessible file. Access and login restrictions are not bypassed.
5. A PDF downloadable from its reference is distinct from source archiving and text extraction. Linking a PDF does not prove external synchronization, extraction or AI analysis. AI analysis is a separate request with its own usage. Stored originals or extracted text do not mean AI analysis has completed.

## Bibliography and PDF visibility

New public-paper bibliography is visible by default to active signed-in platform users, not anonymous visitors. PDF originals, private analysis and research spaces retain their existing access. Reusing an existing private source does not publish it. Edit bibliographic fields and grant/project/researcher metadata on the permanent literature page. External Zotero originals are not overwritten.

## Find and attach PDFs from ChatGPT

Example: “Find references in this space that lack a PDF. Locate public originals, verify the title and DOI, and attach them to the existing references while preserving earlier files.” Tools import a public HTTPS PDF or an actually transferable ChatGPT file and verify bytes, size and hash. They never bypass login or paid access. Original preservation, text preparation and AI analysis are separate states; an unavailable file transfer is not a saved result.

## Compare literature and save to the same work

Example: “Compare these three papers for the Paper Goal and Literature Survey work. State the PDF version and actual coverage; tabulate model structure, data, metrics and limitations. Save the delivered analysis as another request/result pair in this same work and attach the comparison file.” The existing work ID and earlier results are preserved. An excerpt is never labelled a complete-paper reading. Specify explicitly when a new derived work is intended.

## File transfer and retry

Example: “Save the actual image and BibTeX files just created as outputs of this work. Report each saved status and download link.” Attach up to 50 files, combined 100 MB. Reuse the request key for retries; reselect files to renew expired transfer links. The file tool must actually be available in the ChatGPT host. If unavailable, transfer the files through Add sources on the platform. A filename or temporary link is not a preserved file.

## Edit, export and apply requests

Search, select required references or all references, then use the download icon beside search for BibTeX, RIS or CSL JSON. Example: “Correct the selected authors and DOI while preserving the original.” AI summaries run only after a button action or explicit request. Applying the refinement preview updates both the current composer and its local draft, ready to inspect before sending. Dashed graph edges indicate references; solid edges indicate research flow.

## Start paper analysis in the same work

In the existing work composer, use Add sources to select one reference or PDF. Enter a goal-specific question, choose Paper analysis and discovery in the Send to AI menu, and send. Source reading, backward/forward/topic search, registration, synthesis and manuscript drafts are checkpointed. Search uses the public paper title or your explicit public query, not private goals. Verified references and final recommendations are registered automatically; PDF collection is queued with actual storage status. Missing originals limit analysis to the abstract. Forward and similar searches cover the first 25 results each; request further searches separately. Paywalls and institutional login are not bypassed; unavailable originals retain a reason. Registration, original collection, AI analysis and author review are separate states.

## Goals, notifications and collaborative review

Use work menu → Research conditions to edit the question, venue, deadline and constraints. Expand Collaborative guidance to opt into a reviewed goal snapshot and next-action notifications. Only the published goal version is shared; private edits do not silently replace it. Notifications open the same suggestion ID used by the composer and do not execute AI. Deferred or dismissed suggestions are not resurrected. Before joint use, check work sharing, access to originals and the shared goal separately. Approved methods also require use and provider-processing permission. Participants use the research team destination for follow-up analysis. Private results are not published automatically. A researcher confirms submission or goal completion.

## Improve v1 into v2

Use work menu → Know-how settings to select exact approved versions. The default direction and relevant common/activity guidance are pinned; four stages are not required for every work. Existing work-only know-how remains available in collapsed history. Select currently accessible, non-deleted request/result evidence and distinguish paper-specific fixes from reusable improvements. Compare against the actual active v1 using matched input/model/work conditions, held-out input and human review receipts before considering v2 promotion. Promotion changes future defaults and is reversible; it does not rewrite existing results. The built-in method before initialization is distinct from a reviewed and evaluated v1. Method versions, manuscript revisions and model IDs are distinct. Output-only use, source reading and derivation are separate permissions.

## Delete and restore request/result pairs

The author or an authorized space manager can move an accessible request/result pair to trash through its menu. Removed content is excluded from new analysis, method improvement and downloads; other researchers' follow-ups and existing quotations remain. Restore through Deleted requests and results; this never restarts AI. Management does not grant access to another person's private conversation. External copies and provider costs already incurred are unaffected.

## ChatGPT, consultation, Overleaf and costs

Continue in ChatGPT preserves the same work ID and selected evidence, saving only delivered results. Survey request/status/resume/external, method reads/improvements and pair delete/restore reuse consultant\_thread\_prepare/commit reviews. Independent host analysis never claims to have executed a protected method. Core instructions are excluded from ordinary results, history and exports. Provider processing permission is reviewed separately; this does not guarantee that outputs cannot reveal or imply knowledge. Download the draft and BibTeX, review them, then incorporate them in Overleaf; there is no automatic bidirectional sync. Per-request, daily and provider limits pause work while preserving completed stages. Inspect saved stages and usage before explicitly resuming. Team consultation tools currently support requests and status; use the web for resuming and method management. Host tool availability, new server PDF collection, AI quality and Lightsail latency require separate verification.

## Keep only the know-how you choose from a discussion

In the same work composer, ask: "Separate the literature-analysis and manuscript-writing know-how from this discussion. Propose only the changes to each existing method." Only delivered, saved requests and results are evidence. External ChatGPT or Codex can prepare structured proposals using returned evidence IDs and current permissions, then use the existing reviewed save contract. On the proposal page, select only the drafts you want to keep and review the save. Unselected proposals and active versions remain unchanged. Preparing or saving a draft does not validate quality or approve v2.

## Apply several methods together

Use the applied know-how count below the work title, or work menu → Know-how settings, to manage defaults. No empty badge appears for zero items. Search names, handles or activities and review exact versions and roles before saving. Add to a work on the independent detail page lets you choose a space and work and pin an exact version. Type @ in a follow-up, new research request or personal AI consultation to search researchers and know-how together. Choosing know-how adds it only to that request without changing defaults. [[ source citations keep their existing behavior. Multiple selections never auto-merge assets or adopt newer versions. Web, consultation and ChatGPT handoff check the same IDs and current permissions.

## Split, merge, originals and shared evidence

One discussion can produce independent proposals that reference the same segment as shared evidence. Drafts start privately. Separate paper-specific corrections and unverified scientific conclusions from reusable guidance. Change organization prepares a split or merge request in the existing composer for review and sending. A merge proposes a new asset linked to exact parent versions. It does not delete or overwrite the versions, history or active state of assets A and B. Changing how a proposal is organized requires a newly reviewed proposal, not an unreviewed browser-side rewrite.

## Version inheritance and evaluation

Derived work inherits the exact selected versions of independent assets. For existing work-specific methods, the collapsed Work know-how controls retain version and candidate selection in the original work. Review generalized instructions and create a separate asset before reuse in other work. Execution applies modules relevant to the current stage, such as literature, synthesis, writing or review, while retaining links to other stages. New versions are not adopted automatically. Initial approval reviews the actual instructions and their generalization; it does not claim measured quality improvement. Improvement candidates require actual runs with matched questions, inputs, models and effort, held-out cases and human review before adoption. A candidate trial is a one-off selection that does not change the default and uses the normal AI allowance.

## Protection, storage and execution boundaries

An independent detail page shows introduction, ID, handle, authors, usage, current permissions and versions. Editing the introduction or image link does not change execution guidance; image access is checked separately. The default direction is protected guidance. Write a draft or propose a direction improvement, review it and select which drafts to keep. Saving an introduction or candidate does not change the active version. Ordinary catalogs and proposal cards exclude protected instructions. Source review requires current permission and AI-provider processing permission is separate. Reads, selections and structured proposal saves do not run AI, while sending analysis or reorganization requests may use your allowance. Independent external analysis does not prove protected-service execution. Instructions belong in encrypted backups, not code or static files. AI quality, actual ChatGPT tool availability and server speed require separate validation.

## 4. Continue in AI consultation or ChatGPT

Choose Review with AI on a section or download its review context and deliver it to your preferred conversation. Inspect the pinned version, metrics, explanation, source references and omitted evaluation together. Connected consultation rechecks current source access. Reading or downloading does not call AI; actual consultation follows the existing AI budget and permissions.

In ChatGPT or Codex, explicitly provide an accessible result link or the downloaded context and ask: “Check the code provenance and evaluation conditions, identify limitations and propose the next experiment.” Work only through tools actually exposed by that host. A link does not imply that all external code was read or private repository access was granted. Generated code, synthetic checks, real execution and research success remain distinct.

## 5. Cost, limits and recovery

Registering, reading, composing or sending results to Notion does not itself call AI or start a GPU. Server storage/transfer and existing provider charges are separate; zero cost is not guaranteed. Keep models and large datasets on existing Colab, servers or Drive and send only selected small outputs. SDK JSON is limited to 16 MiB, each image to 8 MiB, and tables to 2,000 rows/100 columns. Notion export accepts 40 sections, 10 images and 99 rows/40 columns per table.

On transport failure, the SDK preserves the selected bundle in `.luna-repro/outbox` and returns pending. Treat that folder as private research data. Fix the connection, key or permissions, then explicitly retry with `python -m luna_repro sync`. Reusing a request does not duplicate its result. For authorization errors, check project membership, token expiry and revocation; for HTTP 413, reduce selected content or image size.

The default project quota is 1 GiB of result JSON and image bytes. Retrying the same request consumes no additional storage. HTTP 409 storage errors require reviewing selected outputs or administrator storage policy; HTTP 429 requires waiting before retrying. Default intake limits are 120 requests per user and 600 per project per 15 minutes, separate from AI daily allowance. Earlier results are not automatically deleted to free space.

Notion delivery distinguishes queued, sending and confirmed outcomes. An uncertain write is not automatically repeated. Inspect the destination and connection before recovery; it is not reported as published until confirmed. Platform versions and bounded image bytes are included in existing encrypted operational database backups. Private results are not automatically copied into the lab readable archive.

## 7. Account closure and research evidence

Withdrawal ends login and project access and revokes reproducibility connection keys. Results, immutable versions, images and authorship history are retained. Queued Notion delivery checks current access immediately before execution and stops when access has ended. Copies already sent to external services are not automatically deleted.

Accounts with reproducibility results, composed reports or Notion delivery history require preservation or transfer review before permanent erasure. Erasure does not rewrite the authors of immutable evidence.

Accounts with connection keys but no research history can erase those keys through the existing permanent-erasure process.

## Find reproducibility results and manage connection keys

Find the same materials under All reports and Reproducibility results in the web Reports area, then compose a report from selected fixed-version sections. Settings → Reproducibility connections manages your keys across projects; Project settings → Connections and research tools opens the same manager with that project selected. Participants without project-settings permission can manage their own connections from Settings or Reports. A key belongs to its issuer and one project, and its secret is shown only at issuance. Never put it in consultation, prompts, attachments, source or Git. Issuing and revoking keys are web-management features; do not claim a ChatGPT tool exists unless the host exposes it.

For SDK installation help, state SDK 0.3.0, result contract `luna.repro.page.v1` and your environment. Use Copy connection cell in web Connect results, execute it in Python/Colab and approve in the browser. Finish with `repro.select()` to review local candidates, titles, order and preview before explicitly sending. `repro.publish(*objects)` immediately sends the supplied objects. Default recommendations use no AI, do not infer exact computation cells and cannot recover lost output. In .py, use candidate summaries, choose, preview HTML and publish. The connect confirmation code and approval URL are not a key secret, but never approve an unfamiliar request. Do not claim that a ChatGPT tool operated the notebook runtime or approved a browser connection. Use the hash-pinned command in web Connect results and the public SDK guide <https://github.com/lunalaib-ai/luna-repro-sdk/blob/v0.3.0/docs/installation-en.md> . Public SDK source does not expose research results or keys. The Colab example defaults to local computation; installation or opening a link does not publish. Do not silently upgrade to a new release or request a GPU/paid execution service.

Example: “Check this result's pinned version and code provenance, then identify omitted evaluation conditions and limitations.” Use only results readable under current permissions. Project-management permission does not permit reading or revoking others' keys. Opening lists, menus or connection settings does not run AI, start a GPU, publish to Notion or complete independent reproduction. Notion delivery requires separate destination/content review and an actual remote receipt; synthetic checks are not evidence of real publication.

## Step-by-step practice, syntax and downloadable manuals

Detailed practical guide: <https://lunalaib-ai.com/workspace/repro/guide/?lang=EN>

View the finished example without installing: <https://lunalaib-ai.com/workspace/repro/demo/?lang=EN>

Follow Colab, Windows or Linux setup, the final-cell workflow, incremental sections, text/metrics/table/image/figure syntax, code cells/lines/pinned commits, versions, reports, Notion, AI review, the Pasa adapter and upload recovery. Download standalone HTML/Markdown manuals and runnable `walkthrough.py/walkthrough.ipynb` in the practice ZIP. Read the HTML offline or print it to PDF in your browser.

Inspect the public synthetic example first, then create your own private project demo in Connect results. Code practice saves locally by default; only explicitly choosing to publish uses the connection key. Distinguish a published receipt from a pending local bundle. Neither synthetic example measures research performance or independently verifies reproduction.

Public SDK installation and upgrades:

<https://github.com/lunalaib-ai/luna-repro-sdk/blob/v0.3.0/docs/installation-en.md>

SDK 0.3.0, the `luna.repro.page.v1` contract and the platform version are independent. Public SDK code does not expose keys or research results; installation does not alter existing results, connections or history. Record the installed version and requirements before changing it; test synthetic examples in a new environment. To roll back, reinstall the recorded prior wheel and SHA-256 in that environment while preserving result JSON and outbox files. SDK/CPU practice needs no separate GPU or paid execution service; Colab resources, servers, storage, networking and optional AI have separate costs.

For ChatGPT or Codex questions, specify SDK `luna_repro 0.3.0`, the `luna.repro.page.v1` result-page contract, selected result JSON and error code to avoid confusion with legacy ReproKit events. Never include connection keys in prompts or attachments. Provide review context or an accessible section link, then continue within the actually exposed tools and current permissions.

## 8. Arrange narrative cells freely

In Compose from results, choose Add narrative cell. Use Markdown for the background, methods, data, analysis or limitations where needed. Move result and narrative cells with the up/down buttons or drag their handles on desktop. Add results can bring in sections from other reproducibility pages. Save your changes. This editor does not modify measured values inside result cards.

Saving records a report revision. Open Version history to inspect and restore earlier content; restoration creates a new revision. If another window has edited the report, review the current version first. Removing a result cell only removes it from this report; its source result remains.

## 9. Draft with AI from selected evidence

Save your report, then open AI draft. Select earlier materials or reports that you can currently access in the same lab, and add instructions. Images and saved source code are sent only when their respective options are selected. A link alone does not fetch external repository or Colab code. Review the selected counts, input scope and usage policy before requesting the draft. Reading and previewing do not run AI; generation uses existing AI access, daily allowance and request policy.

AI proposes a title and commentary on background, methods, data, experiments, analysis, limitations and discussion. Unsupported details must remain open questions. Select the proposed narrative cells and title you want to apply, then review them yourself. A changed report or evidence blocks overwriting with an outdated draft. Refine this cell with AI preserves its topic and appropriate length, updating only that draft cell with Undo. It never changes the title, other cells or result order; saving the report stays explicit. Edits made during generation require comparison before applying. Follow real progress and elapsed time for the same job. @mentions and `[[material/report citations]]` work in newly added cells; cited source text uses

current access. This does not perform training, inference or independent verification.

## 10. Move results to Trash and restore

For a result you own, choose → Move to Trash on its card or detail page. Review the impact on referencing reports and ongoing Notion delivery before moving it. It disappears from normal lists, search and reads, while its ID, versions, images and reference history remain. Restore returns the same ID. Management privileges alone do not expose another person's private result.

An SDK retry cannot automatically restore a trashed result. On `result_trashed`, explicitly restore your result in the platform before retrying. Ongoing Notion delivery requires inspection; external copies are not automatically erased. Referencing reports, AI and downloads recheck current source access after revocation or trashing. Permanent erasure of research evidence is not offered.

## Reproducibility report editing, AI drafts and selected code

Use the existing `consultant_thread_prepare(operation=save, values={repro:{operation,values}})` envelope. `report_read(report_id,revision?,offset?)` reads the same web report ID, revision, ordered cells and history. `report_save(title,cells,request_key,report_id?,revision?)` previews a saved result/commentary document. Result cells are `{id:UUID,kind:result,selection:{page_id,revision,section_key}}`; commentary is `{id:UUID,kind:markdown,text,origin:human/ai,references:[record:UUID/report:UUID]}`. Preserve existing cell IDs. Use `revision=0` for new reports and the current returned revision for edits. `ai_prepare(report_id,base_revision,request_key,sources?,images=false,code=false,target?,instruction?,language=ko/en)` reviews one additional paid platform draft and does not generate before commit. `ai_read(run_id)` reads status/draft; `ai_apply(run_id,selected:[proposal cell UUIDs],use_title=true,request_key)` previews adoption. Send the exact review and `preview_token` to `consultant_thread_commit`; `consultant:write` and `report:write` scopes and current domain permissions are checked again. Ordinary repro reads exclude source text. Explicit `source(page_id,revision,section_keys,include_code=true,offset?)` reads only preserved selected code. Never represent a filename/link as code inspection or execution. Save commentary already delivered by the host with `report_save` to avoid a second platform AI call. Verify actual host tool availability; otherwise provide the web editor link and do not claim execution. This contract does not publish to Notion, change sharing or run a GPU.

## External-source sync and originals

Material search and fetch reuse the web library IDs and current access rules. `source_connection_status` reads last check, complete scan and partial collection state for a connection; it never collects or calls AI. Use the web connection preview to map Notion databases and repair past classification. OneNote requires Microsoft read consent and research-section selection on the web. Body collection, attachment preservation and image or handwriting interpretation are separate states. A download button or historical storage receipt is not proof that an original is readable now. Follow the material detail verification and recovery guidance. Server support for a new tool does not prove current ChatGPT host availability.

## Reconnection and current source status

Temporary duplicate refreshes can recover the same token response within a short window. Overall consent currently lasts 30 days; automatic short-lived access token renewal does not extend it. Revocation, overall grant expiry and client changes require sign-in. Check the expiry in My connections. Server support does not prove that the current ChatGPT host is connected. `source_connection_status` distinguishes authentication, access and collection evidence and the need for separate preservation verification, without installation, collection or AI calls. Drive source-read consent is separate from Drive preservation consent. Reconnecting cannot resolve every permission or app configuration problem. Notion classification diagnostics reconcile body retrieval, material registration and report visibility for direct child pages and database entries. Slack requires administrator installation approval and channel membership. Collected materials share their IDs and current access across the web, personal consultation and ChatGPT. Status reads and expanding list summaries use saved state without AI charges.

## Grading integration boundaries

Use the authenticated `/workspace/grading/` web workflow. ChatGPT does not automatically receive answer keys, grades or personal reports and does not grade exams. Exam and snapshot IDs and current permissions in the web service/database remain authoritative. Instructions embedded in documents are not execution authority.

## 39. Share and cite links

Share links or messages and understand recipient access.

Open this guide online

### Use the small share icon

1. Click the small share icon on a report, material, grant or project page.
2. Copy link copies the URL; Copy title and link adds a citation-friendly title. Shared links preserve the current interface language (lang=EN or lang=KO). Paste into KakaoTalk, SMS or Slack and send it yourself.
3. Share to another app opens the device share sheet in supported browsers. Available apps depend on the device and installation. Use link copying when unavailable.
4. Sharing grants no access: the recipient still needs login and scope membership. Document bodies and search filters are not automatically transmitted. Use Share to channel for internal messaging, as described below.

### Share a chat message to another app

1. Press Share on a message to review the sender, content, direct link and LUNA reply guidance. Sharing uses the displayed translation, or the original when viewing the original; Show with original includes both. If the source or access changes, reopen the message and check its translation before sharing.
2. Share to an app opens the supported device share screen. Choose an available app, then its conversation and send action. Available apps and target selection depend on your desktop/mobile device and installed apps.
3. If prompted, press Open app chooser once more. Use Copy text and paste for unsupported browsers or apps absent from the share screen.
4. Closing the dialog clears prepared content. The link grants no access; handing content to the share screen does not confirm delivery.

### Share to a channel with mentions and citations

1. Choose Share to channel and select a channel you may access and message.
2. Explain why you are sharing it. Type @ and select a suggested researcher, or type [] and select a document citation.
3. Review the destination and link, then send. Channel membership does not automatically grant access to the linked source.

4. Verify the link and note in the conversation. If a recipient cannot open the source, check their scope participation and source sharing.

### When mobile app sharing or copying is unavailable

1. Open message sharing and review the channel, sender, time, body and direct link. Share to app opens the device share sheet; choose the recipient app there.

2. Sharing uses the displayed translation or original; Show with original includes both. If the source or access changes, reopen the message and check its translation.

3. Copy text copies that shared text. Check the pasted result in the destination app. Original links still require the recipient's LUNA access.

4. If browser sharing/clipboard support or permission is unavailable, read the fallback guidance. Select and copy from the text box or try a supported browser.

5. Closing/cancelling the share sheet is not successful delivery. Complete sending in the chosen app and check the result. The platform does not automatically choose a recipient conversation.

## 40. Discussions, follows and publications

Discuss research, follow updates and turn decisions into GitHub work.

[Open this guide online](#)

### Comments, likes and tasks

Write comments beneath materials, reports, events, grants and projects; mention people with @username. Follow a comment or discussion for new-comment and edit notifications, and use the same control to unfollow. Settings → My discussions collects your comments, mentions and followed threads. Create a GitHub action item lets a manager choose a related board, review the text, assignees and status, then confirm the remote operation. Current source access and project management authority are required. Popular pages use likes, comments and bookmarks on currently accessible pages.

### Profiles and follows

Enter affiliation and position in Edit profile. Researcher profiles provide follow, unfollow, block and unblock controls. Counts are shown; lists are private by default and visible only to the owner and platform administrators. Owners can make the lists visible to active registered users.

### Digests and email

In Notification preferences choose a digest interval and categories to batch. Other categories stay immediate; leave chat categories unselected to retain instant chat notifications. Inbox details are collapsed and device banners omit body text. Email contains currently authorized details and source links, configured to open a new window; the mail client controls actual window behavior.

### Automatic metadata and connections

Automatic classification associates source projects, grants and identified authors, assignees or connection managers. Calendar attendees match unique account emails within the same lab; association does not prove attendance. Manual changes are preserved. Generate tags extracts keywords locally without sending documents to external AI. Reader changes use the existing proposal and approval process. GitHub repositories and Projects boards are separate connections; organization administrators still need appropriate credentials. Overleaf and OneDrive connections register source links only. An Overleaf GitHub mirror can be imported through a repository connection.

## Publications and AI magazine

Use the connection button in Publications on your own profile to select a collection containing only publications you authored. Provide its library type, numeric ID and collection key; private libraries require a read API key. Daily refreshes are merged by DOI or matching title, authors and year. Lab management → AI research activity magazine enables daily, weekly or monthly features. AI compares source-attributed activity, selects a researcher and writes the article. The selected identity and its evidence IDs are validated; insufficient evidence permits deferring the selection. This is not an objective ranking of effort or research quality. AI permissions, quotas and current source access still apply.

## Read and follow discussions; check GitHub registration

1. Use the discussion composer below materials, reports, events, grants and projects. Existing comments are collapsed and loaded when you open the discussion.
2. Reply, like or follow a comment/thread. Follows can notify you of changes and new replies; unfollow in the same place.
3. Settings → My discussions collects your comments, mentions and followed threads. Check notification preferences for likes and follows too.
4. To turn an instruction into work, a manager chooses Create GitHub work on the comment, reviews the board, assignees, fields and text, then confirms.
5. Check the existing registration status and linked issue first. Additional registration is possible, so verify that another issue is intended.

## Follow researchers, control lists and start a conversation

1. Use follower/following counts and Follow/Unfollow at the top of a researcher profile. Open the adjacent follower/following list control for names.
2. Set list visibility on your own profile. Private lists are visible to you and platform administrators; public lists are visible to active registered users. Following grants no source access.
3. Use block/unblock for someone following you when needed. This is separate from ordinary conversation and source membership permissions.
4. Choose Message on a profile to open a direct conversation. Check the other participant before sending.
5. Overview → Network news → Following shows readable updates with related researcher names. Absence from this feed is not evidence of inactivity.

## GitHub issue comment alerts

1. Link your exact GitHub ID in Profile settings; names and emails are not used to guess accounts.

2. Platform replies notify authorized thread participants/followers, linked issue authors/assignees/participants and users mentioned by platform or GitHub ID. Current read permissions and notification preferences apply; your own reply does not notify you. Follow a discussion to receive replies across collected versions of the same issue.
3. Remote comments notify after the next successful collection. Enable automatic collection or collect now. Signed GitHub webhooks can use `issue_comment` as a collection signal; this is not guaranteed instant messaging.
4. Initial baselines do not flood you with historical comments. Collection uses the latest 50 observed comments and may be limited by remote permissions or collection bounds. Read the original issue for the complete thread.
5. Platform and GitHub comments stay in their original locations.

## Mentions and citations in editors

1. Type `@` in a note, body, comment, review or AI instruction field. Search by part of a username or researcher name, then select a result to insert the exact `@username`.
2. Type `[]` and part of a document title. Typing `[]` alone also lists candidates. Selection inserts a title with its platform link, even if you already typed closing brackets.
3. Use arrow keys to select and Enter or Tab to insert. Enter does not submit while suggestions are open. Escape closes suggestions. During Korean composition, Enter remains an IME action.
4. This covers material/report notes, calendars, discussion edits, chat/sharing, announcements, wikis, magazines, support, Timebank, grant introductions, research questions, biographies and generation instructions. Newly expanded discussions work too.
5. Select a channel before sharing. Only its readers and documents accessible to every reader are suggested. Other editors follow your current access. Inserting a mention does not create notifications in every feature; existing notification rules still apply.
6. Save or submit using the page controls.

## Contributions and profile badges

Open Researchers → Contributions to filter by period and lab and compare up to four people. Profile contributions start collapsed with a level, scores and a featured badge. Expand to see contribution types and next-level requirements. Platform administrators can view and compare profile and directory summaries and activity details regardless of sharing preferences. Contribution summaries are visible to signed-in members by default. To opt out, open Contributions → My sharing settings, uncheck sharing and save. Existing saved private choices are preserved. Summaries may include restricted project activity, while source access stays unchanged. Withdrawing sharing immediately removes public rankings. Scheduled aggregation includes published wiki work, durably stored shared uploads, accepted proposals, shared discussions and confirmed Issue creation. AI chats, likes and visits do not earn points. External PRs, reviews, task completion and support tickets are not yet scored. See the detail page for daily caps and pilot criteria. Report missing or incorrect evidence through Customer support. Platform administrators can exclude or restore readable evidence with an audited reason, applied at the next refresh. Participation levels do not grant permissions or measure overall research ability and are separate from magazine awards.

## Contribution reminders and personal goals

Open Settings → Contribution notifications to receive milestones and weekly summaries, weekly summaries only, or no reminders.

In-app delivery is the default, limited to once a day and twice a week.

The introduction appears once; old badges are not replayed and inactive weeks have no summary.

Email requires both a separate opt-in and your existing email delivery preference.

Choose an optional badge goal or snooze for 30 days.

Expand your profile contribution section for this week across your labs, the change from last week, and remaining points, active days and contribution types for the next level.

No AI allowance is used.

# 41. Favorites, trash and activity review

Organize favorites, restore items and interpret AI reviews.

Open this guide online

## 1. Save favorite pages

1. While signed in, select ☆ near a material, report, grant, project, researcher or conversation title. ★ means saved; select again to remove it.
2. Open Profile → Favorites. Basic search covers titles and accessible material/report content. Advanced search selects grants, projects and researchers; enter tags separated by commas. Choose a type, added dates and ordering as needed.
3. Remove favorite only removes the bookmark. The source stays intact. Revoked pages hide their titles and can still be removed. Favorites are personal settings and are excluded from lab research-data Drive backups.

## 2. Restore or permanently delete

1. Open Profile → Settings → Trash. You can act on items you created or currently manage within your authority. Search by title or type.
2. Restore removes the app deletion flag. Restore a deleted parent grant/project first. Separately deleted children require individual review and restoration. Recheck connection credentials and automatic collection settings.
3. Review permanent deletion shows dependent records and external copies. Enter your current password and confirmation. This database deletion cannot be undone. Original Drive files, historical archives and operational backups are not automatically erased.

## 3. Read AI activity scores

1. Above Participating projects on a researcher profile, view the /100 overall score, four criterion scores, evidence coverage and update date. Profile viewers see this summary.
2. Only the researcher and platform administrators can open Detailed review and criteria for qualitative observations, next steps and weights. Evidence links are filtered by current access.
3. Each 0-4 score is multiplied by 25. A weighted overall score requires at least two assessable criteria. Missing evidence means unassessed, not zero. Do not directly compare scores with different coverage.
4. Platform administrators open Criteria settings from the detail page to edit guidance, automatic refresh and four weights totaling 100. Existing reviews retain their original date and criteria. Eligible reviews refresh through the background runner at the configured interval, subject to AI limits. Administrators can

set the evidence window to 1-365 days (default 30).

## Bulk trash actions and qualitative reviews

Open Profile → Settings → Trash, select authorized items, then Restore or Permanently delete. Select All matching items to include other result pages, up to 500 items. Permanent deletion requires reviewing targets and dependencies, your current password and confirmation. Selected parent grants and projects restore first. External originals and Drive history are retained.

Only the researcher and platform administrators may open detailed AI reviews. Even without numerical scores, reviews provide qualitative observations, limitations, next steps and source links. Older reviews are dated and never treated as current scores. Only platform administrators may edit the rubric and weights.

## 4. Organize favorites into folders

1. Open Favorites from the profile menu. Search is first; the compact folder selector offers all, unfiled and individual folders.
2. Expand folder tools, enter a name and choose a parent for a subfolder, then Create folder.
3. Select card checkboxes, expand Move selected, choose the destination and move. Moving to Unfiled removes folder assignment.
4. Select a folder to change its name/parent. Delete folder only does not delete source documents; move child folders elsewhere first. Bookmarks in a deleted folder move to its parent, or to Unfiled when there is no parent.
5. Cards show essentials such as title, type and folder. The basic list has 18 per page. Advanced searches inspect accessible bodies/metadata and can take longer for larger candidate sets.

## 42. View related GitHub work on a profile

Find related GitHub work and configure deadline reminders.

Open this guide online

### 1. Set your GitHub username

1. Open your profile and find Related GitHub work. If no username is linked, choose My GitHub ID settings on your own profile.
2. Enter the username from `github.com/example-user`, namely `example-user`. Do not enter a repository, email, organization or authentication token.
3. Save and reopen your profile. This settings shortcut appears only on your own profile; another researcher must edit their own ID.
4. Linking an ID neither collects data nor grants GitHub access. Relevant participating scopes need an accessible connection and collected work.

### 2. Find assigned, authored, discussed and directed work

1. Find Related GitHub work below AI assessment. It includes assignments, authorship, comments, participation, mentions, app-created items, linked instructions and explicit researcher metadata.
2. Choose Open, Done or All, then narrow by relation such as Assigned or Authored. Creating an issue does not grant access to someone else's private project.
3. Four items appear per page. Use Previous, Next and the page indicator to replace the work cards. There is no additional hidden remainder within a page.
4. Expand an item's Details to inspect its grant, project, board, priority and due date. Use Open on GitHub for the latest state. A shared issue may show several board memberships.
5. For an empty list, check the username, status filter, active participation and collection. Expand Collection notes for partial or delayed collection warnings.

### 3. Configure assigned-work deadline reminders

1. Use Deadline reminder settings on your own profile, or open Notification preferences.
2. Assigned GitHub task reminders are enabled by default with one-hour and one-day lead times. Change the selections or turn reminders off as needed. Choices range from 15 minutes to one week; previously selected additional lead times are retained.
3. Since Target date has no time, set the deadline reference time in Seoul. For 18:00 with a one-hour reminder, the target is 17:00 that day.

4. Check the category's in-platform, email and browser delivery settings and save. Deadline reminders do not wait for a digest.
5. These reminders cover unfinished work currently assigned to you. Authorship, a mention or a comment alone does not qualify. They depend on collected Target dates and current access, so keep collection healthy.

#### 4. Today's research: plan, check in and share with your PI

1. Open the profile menu → Today's research. Check the date, then review assignments, deadlines and events in Today's direction. Use the date picker or arrows to view other daily records.
2. Choose Write or Edit beside Today's goal, enter your priorities and save. AI plan draft is optional; writing your own plan is fully supported.
3. Expand Select focus tasks to choose up to three. Use Act on a plan item to prepare a comment, message or other next step. The sections below explain applying drafts and executing actions.
4. Check new requests and follow-ups in Reports & actions. A red number counts new requests, not unfinished assignments or executions.
5. View all work opens My work. A written plan and GitHub task state are separate; check the original task state when work is finished.
6. Detailed guide: <https://lunablab-ai.com/workspace/help/daily-research/?lang=EN>

#### 5. From plans to actions and next-day follow-up

1. Act on Today's research opens a side panel beside your plan; Full page opens the dedicated screen. Items in the Follow-ups list open as full pages. The top Follow-ups link returns to the previous filters, page and scroll position. For manual writing, select a section and inspect the linked plan and evidence.
2. Check the method and destination. A cited issue may suggest commenting on that issue. Expand destination content/recent status to confirm the original you intend.
3. Write the content or explicitly request an AI wording draft. Saving a draft prepares privately and does not send.
4. Review the actual account, destination and body in Preview before execution, then approve. Report/event handoffs still require final saving in their forms.
5. Record results and the next check afterwards. Later review drafts can use this action and confirmed observations as evidence.
6. Detailed guide: <https://lunablab-ai.com/workspace/help/research-actions/?lang=EN>

## 43. AI assessment criteria, evidence window and refresh

Interpret role-specific AI reviews and configure their evidence window.

Open this guide online

### 1. Read the assessment and evidence

1. Read the AI assessment card together with its completion time and any historical-snapshot notice.
2. The subject and platform administrators can open the same completed assessment shown on the profile, including observations, next steps and source links. A newer collected version does not hide an assessment when its cited version remains readable.
3. Criteria use 0-4 points converted to a 100-point scale. Insufficient evidence is unscored, not zero. A weighted overall score needs at least two scored criteria.
4. This is a mentoring draft. Compare the documented work and missing evidence, not just scores across people with different data coverage.
5. Expand the card's details for refresh frequency, next attempt and guidance. Scheduled reviews run in turn without repeated profile refreshes.
6. Leadership assessment applies only to platform administrators, active lab PIs/managers and assigned grant PIs (grant owners). Responsible researchers, project owners and delegated managers use contributor criteria unless they independently hold one of those PI roles. Their project creation and management permissions remain unchanged. Its four criteria cover research direction and goals, validation oversight, delegation and mentoring, and risk/resource/follow-up decisions. Contributor assessment covers progress, reproducibility, collaboration and next steps.
7. The subject's explicit assignments determine the perspective, not the requester's role. Leaders need not personally execute every experiment or coding task; team results are not automatically personal achievements. Role and policy changes trigger a scheduled reassessment. A previous score with a different perspective is hidden until reassessment finishes. An assessment labeled as predating role-specific criteria retains its original interpretation.

### 2. Administrator: set the period and frequency

1. Open Lab management → AI assessment policy, or use Assessment settings from a detailed review. Only platform administrators can save the policy.
2. Enable automatic profile review. Choose a refresh interval of 12, 24 or 48 hours, or seven days.

3. Enter an integer from 1 to 365 under Evidence window (days); the default is 30. Seven includes today and the preceding six days. A 24-hour refresh and a 90-day window reviews the latest 90 days each day.
4. Background goals and plans may fall outside the window. A longer window does not acquire unconnected or uncollected data, and input limits remain.
5. Add guidance and four weights. Each must be 1-100 and the total must be 100. Save, then verify the period and interval.
6. Changes apply to subsequent reviews. The previous completed review retains its original criteria while the replacement is queued. AI permission and the administrator's daily allowance still apply.

### 3. When a review is missing or old

1. Check linked GitHub/Notion identities, participation and researcher metadata on collected sources. Insufficient linked evidence may leave a review pending or partially unscored.
2. Ask an administrator to confirm lab AI and automatic review settings. Queued and running jobs are processed in turn.
3. Temporary failures retry up to twice after an hour; daily quota limits resume after the Seoul date changes. Include the displayed status when requesting support.
4. A readable completed snapshot may remain visible during refresh. Check its date and source set instead of treating it as the new period's result.

## 44. Personal Zotero publications and the lab bibliography

Connect a collection of your own authored publications.

Open this guide online

### 1. Prepare a publication collection

1. Create a Zotero collection containing only publications you authored and sync it online. Keep it separate from a general reading collection.
2. In web Zotero, identify personal versus group library, its numeric ID and the collection key. The collection name is not its key.
3. Private libraries need a read API key with access to that library. Enter it only in LUNA's secret field. Expand the publication connection guide on the settings page for ID examples.

### 2. Connect it to your profile

1. Use the connection button in Publications on your own profile, or open <https://lunalaab-ai.com/workspace/publications/settings/>
2. Enter library type, numeric ID and publication collection key. Enter an API key only when replacing it.
3. Confirm that the collection contains your authored publications and enable display on your profile and in lab publications, then save.
4. Check connection status and the daily refresh result. Review your profile list and the lab bibliography on Overview. Saving does not guarantee immediate collection.
5. DOI or normalized title/author/year matches merge duplicates. Check those fields in Zotero if a paper is duplicated or editions are merged unexpectedly. Disable the display option to stop showing the collection.

## 45. Read and publish the AI activity magazine

Read and publish bilingual research stories and inspect references.

Open this guide online

### 1. Read a summary, full article and archive

1. The People shaping our lab section on Overview shows one latest readable article summary, without inline document citations.
2. Open the full article, evidence and discussion to read the complete story including its summary. Evidence is collected in Sources used in this report below the article.
3. Follow the other articles/archive link to browse previous editions. One Overview card does not mean older articles were deleted.
4. Unreadable or trashed articles may be hidden. An unavailable-source notice means the source's current state cannot be verified.

### 2. Administrator: configure publication and badges

1. Open Lab management → AI research activity magazine for the lab.
2. Enable it and select daily, weekly and/or monthly publication. Set minimum evidence per candidate, selection criteria and article instructions, then save.
3. Set the profile celebration duration from 1 to 3650 days. A badge and article link appear for that period after selection; the article remains archived after the badge expires.
4. Give concrete instructions, such as emphasizing reproducible validation and collaboration, including the researcher's name in a descriptive headline and using an editorial magazine tone.
5. Candidates need collected author/assignee evidence. Selection is not fixed to a named person, but a one-person candidate pool limits comparison. It is not an objective ranking of actual effort.

### 3. Publish an edition now

1. Save magazine settings, then choose a period under Publish magazine now: today, this week to date or this month to date.
2. Press Publish now once. An already running edition for that period opens instead of creating a duplicate running job.
3. Check the linked report's pending, running and completed states. Publication follows AI writing and evidence validation. Requesting a completed period again may create another edition.

4. If candidates lack evidence, check GitHub/Notion attribution and collection within the period. For disabled AI, limits or failure, inspect the settings and job status.
5. Choose selection notifications in preferences. Lumen awards are separate administrator-reviewed Timebank transactions, not automatic payments.

### Insert files and images inside the body

1. Open a new wiki page, Edit page, a new/edit proposal, Write report, or a manual magazine article.
2. Select the Notion connection for a report or the lab for an article first. That lab needs a configured private Google Drive archive.
3. Click the insertion point inside the body. Choose Attach file for PDF, DOCX, HWP, TXT, Markdown, ZIP and other supported files, or Add image for PNG, JPEG or WebP.
4. Limits are 20MB per file, 1.5MB per image and 20 attachments per document. Once Drive storage is confirmed, a download link or image Markdown appears at the chosen position. Images are normalized to safe JPEGs.
5. Check Preview, then save/publish or submit the proposal. Use Retry attachment or Cancel attachment if transfer fails. If Google archive authorization has expired, keep the editor open, use Open archive account settings in a new tab, reconnect the existing account and retry. No new folder is needed. Ask your lab administrator to reconnect if you cannot manage settings. Finish uploads before saving.
6. Before saving, only the uploader can access new files. Pending proposal files are available only to the author and reviewers. Saved files follow the document audience, including Only me and revoked memberships.
7. Notion exports and wiki ZIPs retain authenticated platform links for new attachments. They do not make files public or copy large originals into the ZIP. Download originals using their links. Read-only collected documents are still edited in their original service.

### Korean and English stories for all members

1. Read the latest story on Overview or open the full Magazine archive. Distribution editions are available to all active signed-in members regardless of grant or project participation, not anonymous internet visitors.
2. Choose Korean or English for the title, summary and body. Share, export and favorite icons are available.
3. New or edited articles receive a background bilingual rewrite and separate editorial review. Unpublished results, contracts, personal information, internal links and commit IDs are omitted; stories explain the research question, central idea, current development, its significance and next verification steps in accessible prose. Articles reduced to generic praise or collaboration updates are held at editorial review.

4. Unreviewed originals are not exposed to nonparticipants. Author permissions, AI quota and lab settings may delay preparation. Drafts and Only me articles are not distributed.
5. Original drafts, evidence, attachments and historic internal discussions retain their permissions. Distribution editions have separate discussions. Temporary source outages do not hide an already reviewed edition.
6. Editing, making the original private or deleting it hides the edition until the updated source passes review.

## Researcher profiles and article references

Below the article, select a featured researcher name to open their profile. References and supporting materials list document titles and links only when you can currently read both the source report and each document. The article body has no inline citation links. An empty list does not mean there are no sources; access or connection status may limit availability. Printing and exports follow the same permissions.

## 46. Write announcements and request AI drafts

Write, translate, review and publish bilingual announcements; link existing editions.

[Open this guide online](#)

### 1. Read and write bilingual announcements

1. Announcements and Overview show only editions published in the current interface language. Switching Korean/English opens the corresponding edition of the same announcement. An unpublished language is not replaced with the other language.
2. Platform administrators: Announcements → Write. Enter both titles and bodies using Korean/English tabs or Compare side by side. Save both drafts together; published content stays unchanged.
3. Preview each language's tables, images and links, then select Korean reviewed and English reviewed. Publish both languages makes one announcement available in both language lists.
4. For urgent notices, expand Publish one language first or unpublish and choose Korean only or English only. Review and publish the other edition later.
5. Changing one language clears review of the other. Save the draft first, then compare and review both again. Save before refreshing or leaving the editor.
6. Administrators use Published notices, Drafts and Trash. Drafts show each edition's writing and review state. Existing URLs and access rules remain valid. Check that broad-audience notices do not copy restricted research sources.

### Attach files, images, mentions and document citations

1. Open Announcements → Write, or Edit and publish on an existing announcement. Set the title, font and size, then use the toolbar directly above the body.
2. Choose Attach file for supported PDF, DOCX, HWP/HWPX, TXT/MD, CSV, PPTX, XLSX, ZIP, image, audio or video files. Limit: 20MB per file and 20 files per announcement. Executable and HTML files are rejected. Add files one at a time.
3. If you manage multiple labs, first choose Attachment storage lab. Its Google Drive archive account and folder must be configured. Wait for confirmation and the download link in the body. Retry or cancel a failed upload. Saving and publishing wait until attachment completion.
4. Choose Image to display a PNG, JPEG or WebP up to 1.5MB in the text. Images are resized and re-encoded for display. Use Attach file to distribute the original file.
5. Choose @ Mention or type @ followed by part of a username or name and select a result. Choose Cite a document or type [[ and part of a title to select an accessible document. Arrow keys and Enter/Tab also

work. Limit: 12 citations.

6. Preview images, tables, citations and attachment links. Saving a draft sends no alerts; publishing follows announcement notification preferences. Only announcement readers can download attachments. Remove the link from both published editions and publish again to revoke reader access. Saving a draft alone keeps the live attachment; this does not delete its Drive original.

7. Attachments are for sharing files; they are not automatically analyzed or executed.

## Check unread announcements

1. Open the profile menu. The number beside Announcements counts published notices you have not opened. It disappears when all notices are read.

2. Open a notice from the list to mark it as read. Viewing a list or notification, or downloading an attachment alone, does not mark the notice as read.

3. Counts refresh when opening the menu, returning to the tab and approximately every 15 seconds while the page is visible. Drafts and unpublished notices are excluded. Editing a previously read notice does not make it unread again.

4. Existing view records are used; visits before read tracking was introduced cannot be inferred.

## Administrators: check announcement readers

1. As a platform administrator, open an announcement and choose Read activity at the top. Researchers, grant leads and lab managers cannot access the list unless they also have the platform administrator role.

2. Review the reader count, names, usernames and first/last viewed times. Use Previous and Next for pages of 20 readers.

3. Current revision viewed means the user opened the current version. Earlier revision viewed means the announcement has changed since their last visit. Reopening it updates this status.

4. Only successfully opening a published announcement detail page is recorded. Lists, notifications, drafts, attachment-only downloads and the reader list do not mark the announcement as viewed. Repeat visits update the last-viewed time without increasing the reader count.

5. Tracking begins when this feature is introduced; prior visits are not inferred. A visit is not an explicit acknowledgment or approval. Permanently deleting an account also removes its read receipts.

## 2. Request an AI draft with /notice

1. Open AI consultation as a platform administrator with AI permission.

2. Start the question with /notice and specify purpose, audience, dates and required content. For example: /notice Draft next week's seminar announcement; mark the undecided room as TBD.

3. Bring the generated title and body into the announcement editor and review facts, dates, audience and citations.
4. Edit, preview and save to publish. Generating a draft is separate from final publication. The AI does not finalize undecided venues, schedules or decisions.

## Share, export and bookmark announcements

1. Open an announcement and use the share, export and star icons above its title.
2. Choose link copy, title-and-link copy, channel sharing or another app. Sharing a link does not grant access.
3. Choose print view, PDF, Markdown or text. Administrator edit/readership links are excluded; download attachments separately.
4. Click the star to add a private favorite and click again to remove it. Withdrawn announcements and revoked access also hide their content and title in favorites.

## Receive new announcement alerts

1. Publishing an announcement alerts active users who can read it. Platform alerts are immediate by default, even without an @mention.
2. In Notification preferences → Notification categories → Announcements, requests and operations, expand the group and choose platform and email delivery. Email and device push follow channel permissions and administrator policy.
3. Select Digest and an interval for grouped delivery. Existing opt-outs and digest choices remain in effect.
4. Drafts do not notify. Editing or republishing the same announcement does not send duplicate alerts. Withdrawn announcements and revoked access do not expose the body through notifications.

## Generate, review and apply AI translations

1. Expand AI translation draft and choose the lab AI settings. Choose Generate English from Korean or Generate Korean from English. Current inputs are saved as drafts first.
2. A platform administrator and enabled lab AI settings are required. A dedicated queue avoids waiting behind document analysis and reports; short titles and bodies are translated together. Longer content is split and completed segments are shown. The requester's personal daily allowance applies, including failed execution attempts. Shared lab analysis attempts are not used.
3. When ready, the translated title and body appear in the target editor without refreshing, and side-by-side comparison opens. Table structure, images, attachment links, numbers, times, URLs, mentions and document references are checked for preservation. Review meaning, names and obligations, then save the draft. Nothing is published automatically.

4. If you edit the target language while waiting, your input is preserved and the candidate appears separately. Compare it before choosing Replace the target draft with this translation. Generate again if the source changed. Mark the draft reviewed and publish to make it visible. Results with changed permissions or settings cannot be applied.
5. Failed or limited requests preserve the original. Write manually or check settings and allowance before retrying. File contents and text inside images are not translated.

## Link existing editions and understand language records

1. Organize legacy notices → Link existing Korean and English announcements appears only when separate published editions remain. Select the two versions of the same notice. Similar posting times are suggestions only; compare the titles and bodies yourself. Linking is never automatic.
2. Preview link → confirm both texts are the same announcement → Link the confirmed announcements. The Korean URL becomes canonical; the old English URL redirects to it. Originals, attachments and reading history remain, with no new alert.
3. Unread counts include notices published in the current language. Opening either edition marks the shared announcement read. Administrator read activity separately records the actual language and published revision. A page visit is not consent or approval.
4. Email/push notify once per announcement when an edition exists in the preference saved by the header language selector. Existing accounts initially use Korean. In-app content follows the current interface language. Added translations, edits and republication do not duplicate a previous alert.

## Trash, restore and permanently delete notices

1. Platform administrators choose Move to trash from a notice list, detail or editor. Review the Korean and English editions and their linked addresses before confirming.
2. Trash retains content, attachments and read receipts while hiding the notice from readers and shared links. Restore returns its previous published or draft state without new notifications or translation.
3. Permanent deletion is available only from Trash. After confirmation it removes both editions, app images, attachment links and read receipts and cannot be undone. Original Google Drive files remain.
4. If another administrator changes the notice, review it again. Restore a trashed notice before editing it.

## 47. Timebank: deposits, transfers and GPU usage

Request credit, transfer lumens and record actual GPU usage.

Open this guide online

### 1. Understand lumens (LM) and GPU time

1. Open Timebank in the top menu. Your available balance is in lumens (LM), shared across all GPU types.
2. One A6000 GPU for one hour equals 1 LM. H100 does not need a separate currency; its configured rate converts hours into the same LM balance.
3. Example: if the administrator sets H100 to 3.57 LM per hour, one H100 for two hours costs 7.14 LM. Two GPUs for two hours each means a quantity of four GPU-hours, costing 14.28 LM. The rate 3.57 is illustrative; use the rate displayed in your account.
4. Available hours on each GPU card equal your balance divided by that GPU's rate. Do not add the hours across cards: they are alternative views of the same balance.
5. This is an internal research-time ledger. Recording an entry does not pay Runpod, start or stop a GPU, or automatically measure usage. Operate and stop the actual server separately.

### 2. Check your balance and transactions

1. Check your available balance at the top. A new account may remain empty until credited.
2. In My transactions, inspect the date, type, reason, sender/recipient, LM change and the GPU, hours and rate used at the time.
3. A plus sign is a credit; a minus sign is a transfer out or recorded GPU usage. Entries are newest first, with 10 per page. Expand an entry to read the full reason, exact amount and transaction time.
4. Submitting a credit request does not increase your balance. Verify the credit entry after an administrator approves and transfers the funds.
5. Rate changes do not rewrite historical transaction amounts. If a calculation seems unexpected, expand the entry and use its saved rate rather than today's rate.

### 3. Researcher: request additional credit

1. Under My account transactions, choose Request credit from administrator.
2. Select the GPU type and check its displayed LM-per-hour rate.
3. Enter the total number of GPU-hours needed. For one A6000 for five hours, enter 5. Do not enter an LM amount in the hours field.

4. Explain the project, planned experiment and time estimate under Reason, for example five comparison runs of about one hour each. Review the calculated amount and save.
5. Credit requests starts collapsed, showing pending and processed counts. Expand it to filter Pending, Processed or All, with five requests per page. Open an individual request to read its full reason and review note. An administrator may approve a different number of hours. Check the actual credited amount after approval.
6. Ordinary researchers cannot top up their own accounts. For an urgent request, identify the existing request to the administrator instead of submitting duplicates.

#### 4. Researcher: record actual GPU usage

1. Confirm the GPU type and time actually used on the server, then open Timebank → Usage record.
2. Select the GPU and enter the hours: 30 minutes is 0.5; 90 minutes is 1.5. For multiple GPUs, add GPU count multiplied by hours.
3. Describe the research or experiment, review the hours and converted LM amount, then save.
4. The amount is deducted from your balance and recorded in My transactions. Separately confirm that the actual server has stopped.
5. Insufficient balance prevents the transaction. Request credit and record the usage after it arrives. Report an incorrect entry to the administrator or Support with its details; do not record the same usage twice.

#### 5. All researchers: transfer to a colleague

1. Choose Withdraw / deposit to another researcher. Active researchers as well as administrators can transfer within their own balance.
2. Select the recipient by both name and @username. You cannot transfer to yourself.
3. Choose a GPU and hour quantity and enter a reason. To provide one H100 hour, choose H100 and 1; at a rate of 3.57 this transfers 3.57 LM.
4. Recheck the recipient and amount before saving. One transaction debits your account and credits theirs.
5. Verify the recipient and debit in your ledger; the recipient checks their credit entry. Credit notifications follow their delivery preferences. For an incorrect transfer, discuss correction with the administrator and recipient using the saved transaction details.

#### 6. Administrator: set up the bank

1. Open Timebank as a platform administrator and locate Administrator bank operations. Ordinary researchers do not see these controls.

2. Register GPU types under GPU conversion rates. A6000 remains the base at 1 LM per hour. Enter the lab-approved conversion for additional GPUs such as H100.
3. Check the time represented by actual Runpod prepayment or budget. Open Administrator account deposit, enter the GPU, hours and supporting reason, and save.
4. Verify the increase in your administrator balance. This credits the internal ledger; it does not make a Runpod payment.
5. Distribute credit using Withdraw / deposit to another researcher. Expand Researcher accounts and choose Deposit beside a person to preselect the recipient. The administrator's balance decreases by the amount distributed.

## 7. Administrator: review a credit request

1. Open the new request notification or choose Review / credit beside a pending request in Timebank.
2. Check the applicant, GPU, requested hours and purpose. Choose Approve credit or Reject.
3. For approval, confirm Approved hour quantity. Keep the requested hours or enter a different amount. The GPU type stays the one requested.
4. Add a review note and save. The approved quantity is converted at the GPU's current rate and transferred from the administrator to the applicant. Sufficient administrator balance is required.
5. Rejection records the decision and note without transferring credit. Check the request list for the outcome. Reviewed requests cannot be approved again. Restore a retired GPU before approving a pending request that uses it.

## 8. Administrator: edit, retire and restore GPU rates

1. Open GPU conversion rates. For a new GPU, enter its name and lumens per GPU-hour, enable Available, and save.
2. Use Edit beside an existing GPU. Enter 3.57 if one H100 hour should equal 3.57 A6000 hours. The new rate applies to subsequent transactions.
3. A GPU with transaction or request history cannot be renamed. Add a new GPU if a different name is needed. A6000's name, base rate and availability are fixed.
4. Retire an unused GPU through its delete confirmation. This removes it from new-transaction choices while preserving old entries, requests and balances. A6000 cannot be retired.
5. To restore it, edit the retired setting, enable Available and save. If another administrator changed the setting, reopen the form and review the latest values first.

## 9. Administrator: award research credit

1. Choose Award credit, or use the administrator-only award link on a featured researcher's profile.
2. Check the recipient, GPU and hours. State the research activity, article or review supporting the award in the reason.
3. Review the converted amount and administrator balance, then save. Credit moves to the recipient and is recorded as an award.
4. Magazine selection or an AI score does not automatically pay an award. An administrator reviews and confirms the recipient and amount.

## 10. Notifications and common problems

1. Open Profile menu → Settings → Notification preferences and find Time bank deposits, withdrawals, usage and requests. Choose in-platform, email or browser delivery and save. Email and browser also require their overall channel switches.
2. Credit notifies the recipient; a new request notifies platform administrators. Administrators can also receive changes to other researchers' accounts, including transfers, usage and awards.
3. For immediate delivery, exclude Timebank from digest categories. Administrator policy or disabled delivery channels can prevent delivery.
4. Missing GPU: ask for a rate to be registered. Insufficient balance: request credit. Missing recipient: check active account status. Unexpected amount: check GPU, hours and the transaction's saved rate.
5. If a save seems slow, inspect the ledger or request list before creating another transaction. If unresolved, send Support the page URL, time of action, expected amount and observed result.

## Read compact transactions and credit requests

1. Scan short transaction summaries and expand a transaction for its reason and approval explanation. Collapsing a summary does not delete the original text.
2. Transactions are paginated. Check date, type, amount and the historical conversion in details; do not confuse it with today's conversion rate.
3. Credit request status starts collapsed. Open its summary, filter Pending/Processed/All and use pagination.
4. Saving or reviewing a request may open that specific item. Check approval, explanation and actual credit state.
5. Collapsing or filtering does not change transactions, balances or approvals. Actual transfers/approvals use their separate confirmation flow.

## 48. Profile and notifications

Choose immediate or digest alerts and troubleshoot device push.

Open this guide online

### Short notification previews and direct originals

Expanded notifications show a short preview of at most 280 characters and four visible lines. Use Open full page above the preview for all details. GitHub issue materials show Original GitHub issue first, opening that issue directly in a new tab. Other materials offer View original when a valid source link is available. Links require current material access. Notifications, originals and translation copies are preserved. No additional AI execution or usage charge is involved.

### Recent material changes and historical mentions

External material alerts require a verified source change within the preceding 24 hours. Re-imports, copied annotations, automatic classification and new snapshots of unchanged content do not revive historical mentions. Undated or future-dated sources do not produce new external material alerts. A new local memo mention on an old material still follows current access and notification preferences. Existing materials and notification history are preserved. This rule starts no additional AI or collection work and incurs no additional usage.

### 1. Profile details and photo

1. Open the profile icon → View profile → Edit my profile.
2. Enter your name, email, phone, website and biography. Prefer an https:// website address.
3. Upload a JPG, PNG or WebP under Profile photo on the same screen. Maximum size is 1.5 MB and 16 megapixels; the saved copy is resized to 512px and location metadata is removed. No photo URL is needed.
4. Save, then check both profile and navigation images. To remove the photo, select Remove current photo and save.
5. New usernames must begin with an ASCII letter and contain only letters and digits. They are stored in lowercase; sign-in is case-insensitive. Existing usernames and URLs remain valid.

### 2. Optional Notion user UUID

1. Expand External account information. GitHub username and Notion user UUID optionally attribute source authorship to your research activity. They are not authentication tokens.
2. Leave the UUID blank if unknown and ask the connection administrator. Do not substitute a page ID from a page URL.

3. An administrator with permission to read user information can use the official Users API (<https://developers.notion.com/reference/get-users>). Find your type=person entry in results and copy only its id. Follow next\_cursor when has\_more is true.

4. Do not enter a bot or another person's ID. Never paste tokens or the entire user directory into conversations or public documents.

### 3. Choose notification delivery

1. Profile → Notification preferences. Platform categories default to all selected. Choose chat mentions, replies, messages, mentions in materials/events, researcher assignments and participating scope updates separately. Administrator-disabled categories cannot be received.

2. Enter followed tags and press comma or Enter. Enable the followed-tag category too. New or updated readable material with these tags triggers an alert; tags never grant access.

3. Enable email and choose email categories → Save. The administrator must connect a Gmail sender. Delivery uses your profile email and includes currently readable author/title/body (up to 12,000 characters) plus direct links. Use an email account you control; open LUNA for longer content.

4. Enable browser notifications → Allow in the browser prompt. Successful device registration changes the button to Allowed and disables it. Enable PC/mobile push → Save. If permission is already allowed but Connect device push remains available, press it to complete registration.

5. Ask another participant to enter @yourusername followed by a space and a message in your channel. Check the bell and selected delivery methods. Explicit self-mentions can also test alerts; ordinary messages to yourself are not notified. Check both personal categories and administrator policy.

6. On Windows: Settings → System → Notifications → enable your browser; check Do not disturb/Focus settings. Allow lunalab-ai.com notifications in Chrome/Edge site settings. Alerts can arrive while using other tabs or programs. Delivery after completely quitting the browser depends on its background settings and OS policy.

7. Android: open the site in a supported browser such as Chrome and allow notifications. iPhone/iPad (16.4+): Safari → Share → Add to Home Screen → open the LUNA icon → sign in → Notification preferences → Allow. Ordinary Safari tabs may not offer push registration.

8. No paid SMS service or app purchase is needed. Network, power saving, focus mode and expired login can limit delivery. The platform bell retains notification history. Connect each browser/device separately; after switching accounts on a device, reconnect device push.

### 4. Administrator: connect a Gmail sender

1. Open Profile → Lab management → Notification policy and Gmail sender. Only platform administrators can access it.

2. Enable 2-Step Verification for the sending Google account. Create an app password for LUNA Lab at <https://myaccount.google.com/apppasswords> . Institutional policies or Advanced Protection may prevent this. Official guide: <https://support.google.com/accounts/answer/185833>
3. Enter the sender address and 16-character app password only in this settings form; enable Gmail sending. Do not use the ordinary Google password.
4. Choose Save and send a verification email to the sender, then check that Gmail inbox. For failure, check the app password, 2-Step Verification and account policies.
5. Select allowed notification categories and save. Users must separately opt in to email. Automated sending is limited to 100 attempts per day overall and 20 per recipient. Old notifications are not sent in bulk; failed or uncertain attempts are not automatically retried.
6. Changing your Google password may revoke app passwords; save a new one if needed. To disconnect, disable Gmail sending and revoke its app password in Google. No paid mail or SMS subscription is installed.

## 5. Choose chat notification events

1. Open Profile → Notification preferences.
2. Select new channels, new messages, mentions, replies, changes and AI response completion separately in the In-platform column. All are initially enabled.
3. For email, enable Receive email notifications, choose categories in the Email column and Save. You can keep all in-platform events but email only mentions and replies.
4. Each message uses one category, in priority order: mention, reply, general message. Disabling that category suppresses it. Ordinary messages you send do not notify you, but explicit self-mentions do. Creating a channel notifies current participants.
5. Open a bell notification to reach the message or conversation. Email includes currently authorized notification content and a direct link. Sending requires the configured sender and is subject to daily limits. Old notifications are not sent retroactively.
6. Losing channel access stops new notifications. Only the owner receives private AI response notifications.

## 6. Clear your inbox

1. Open Notifications, select the checkboxes beside items and choose Delete selected. Select this page selects the currently displayed items, up to 20.
2. Clear all notifications removes existing notifications across all pages after confirmation. Notifications arriving after you opened this screen are preserved.

3. Cleared notifications disappear from your inbox, bell count and direct links; pending email deliveries are cancelled. Mail already sending or sent cannot be recalled. Linked messages, materials and other recipients' notifications remain intact. Server processing history is retained to prevent duplicate delivery.

4. While the site is open, the unread bell count refreshes normally every 30 seconds while the page is active, and after sending a conversation message.

## 7. When a Windows banner is missing

1. In Notification preferences, enable PC/mobile push and desired categories → Save.

2. Press Connect / check device push and confirm Connected. This confirms device/provider registration, not that you have seen an OS banner.

3. Press Test notification. It requests a banner on this device and sends no message or email to anyone else.

4. If invisible: Windows Settings → System → Notifications → Google Chrome or Microsoft Edge → enable notification banners and notification center. Turn off Do not disturb and Focus, and allow notifications for lunalab-ai.com in browser site settings.

5. Verify with a new mention. An open LUNA page also shows a bottom-right notification card and checks periodically while active. Errors or rate limits increase the retry interval, so immediate display is not guaranteed. Provider acceptance does not prove Windows displayed a banner. LUNA cannot override OS focus or power-saving restrictions.

## 8. Connection state after testing

Opening notification settings verifies the existing device subscription and displays Connected. Finishing a test, a temporary send failure or switching windows does not reset the connection. A test displays one notification rather than a burst of existing alerts. Only an actual provider expiration response removes that subscription and prompts reconnection. Device receipt and an OS display request do not prove that a banner was visible. If a Windows banner is missing, follow the device guidance to check browser banners, notification center, Do not disturb and Focus.

## 15. Read notifications without leaving your page

Select the bell to open the notification panel on the right. On mobile it fills the screen; Close or browser Back returns to your page. Choose Last 7 days or Unread, then select a title to expand its details. Opening the panel alone never marks anything read; only a successfully displayed detail marks that item read. In Unread, a just-read item stays until the next refresh to avoid jumping. Last 7 days covers the preceding 168 hours; Unread has no date limit. The actions menu marks unread notifications from all dates through the confirmation snapshot; later arrivals remain unread. Use the trash icon to delete a read notification and Undo within 10 seconds. Delete all read notifications includes older dates and preserves unread notifications. Only your inbox changes, not source material or other recipients. Use the new-notifications button to refresh. Open related page takes you to the source; View all notifications opens the full history and deletion controls. Only one help, conversation or notification panel is shown at a time.

### Notification language and free translation

System notification labels follow the interface language. Translate source titles and excerpts for free using the translation icon and switch back to the original. Translation preferences also apply to research messages. Mobile reuses your saved copies; for new mobile translation use Chrome menu → Translate. No AI allowance is consumed and translation lookup alone does not mark a notification read.

### Preventing repeated external-material alerts

External-material alerts require a source change within the last 24 hours. The same GitHub issue in one lab shares its observation history across repository and board connections. Older changes or comments, collection timestamp labels and automatic tag or assignee maintenance do not create new alerts. A displayed comment author is not a mention. Actual new issue-body mentions and recent new or edited comments remain subject to current access. Legacy snapshots without body-mention provenance establish a silent baseline first. New explicit local notes and existing notification history are preserved.

### Enable device notifications for conversations

1. An automatic guidance card appears on signed-in pages when device notification permission is missing. Click Enable notifications, then Allow in Chrome. If Chrome uses a quiet prompt, check the notification icon in the address bar.
2. After permission is granted, the app connects this device and enables your device notification preference. Existing category opt-outs and digest choices remain unchanged. For immediate conversations, enable Platform and clear Digest in Chat and comments, then save.
3. If already blocked, the site cannot force another permission prompt. Open site information to the left of the Chrome address bar → Site settings → Notifications and allow them, then return. Use Connect device push in notification preferences if needed.
4. Later dismisses the card for the current tab session. Notification preferences always offers the relevant guidance. On unsupported iPhone/iPad browsers, add the site to your Home Screen with Safari and allow notifications in the web app.

5. Permission alone does not guarantee an OS banner. Use the device test button and check OS notification banners, Do not disturb and conversation notification settings.

## 8. Separate digests from immediate alerts

1. Choose delivery channels per category in Notification preferences. Expand a group, change the needed categories and save.
2. Select a digest interval of 15 minutes, one, three or twelve hours, or daily. Immediate delivery does not wait for a digest.
3. Select categories to batch, such as source changes. Leave time-sensitive chat categories unselected so they remain immediate.
4. Configure GitHub deadline lead times and Seoul reference time separately. These use collected due dates for unfinished work assigned to you.
5. Expand administrator-restricted categories and device/Windows guidance only when needed. The count refers to categories. Required in-app support alerts are exempt, and administrators can follow the policy link to adjust restrictions. Test and device-permission/connection controls remain separately available.

## 9. Choose social, Timebank and approval alerts

1. Enable researcher-follow alerts for new followers and magazine-selection alerts when you are featured.
2. Enable discussion likes and comment/thread follow categories to receive reactions to your discussions.
3. Timebank covers credits to your account, administrator credit requests and changes to other researchers' accounts. Follow the Timebank chapter for transactions.
4. Join and metadata approval requests reach currently authorized reviewers. New support inquiries alert administrators and resolution alerts the author immediately in-platform; email/push still follow channel consent.
5. If nothing arrives, check overall channel switches, digest selection, administrator policy and current access. Enabling a category does not replay every historical event.

## 10. Select, mark read and inspect details

1. Select several notifications or select the current page.
2. Use Mark selected read or Delete selected. Source documents, comments and tasks are not deleted.
3. Expand details only for the item you need and use its source links. Device banners omit long details.
4. Email may include readable details and navigation links. Actual new-window behavior depends on the mail client and browser.

## Compact deadline reminder selection

1. In notification preferences, open GitHub deadline reminders and click the reminder-time box. It normally shows only your selected times.
2. Check or clear multiple times. The summary updates immediately. Click again or press Escape to close it.
3. Save at the bottom of the page.

## Test device push

1. Open notification settings on the PC or mobile device you will use and connect device push. On iPhone/iPad, open the Home Screen web app added through Safari.
2. Send a test notification. The server targets this device only and waits about 15 seconds for receipt and an OS display request. This does not prove a banner was visible.
3. If receipt succeeds without a banner, expand device help. Check Chrome banners and notification center in Windows Settings, including automatic Do not disturb rules. On Android, check browser/site notifications. On iPhone/iPad, check web-app notifications and Focus.
4. Test local display only bypasses the server. If the subscription expired, test again to reconnect. Wait 20 seconds between tests.
5. For immediate chat alerts, enable device push and chat categories, clear their Digest boxes and save.

## Delayed collection and source-change digests

1. First collection time is different from the source's actual creation/modification time. External material older than the previous digest is excluded from new-source digests.
2. Sources without dates retain observation-based behavior and may include older material. Missing dates do not prove freshness.
3. No empty digest is sent when no new items qualify. Existing notification history remains.
4. Inspect accessible items in Notifications independently of push/browser status. Polling backs off after temporary errors, so updates may be delayed.

## 49. Drive preservation and research context

Check preservation status and distinguish wiki originals from backups.

Open this guide online

### 1. Reading versus preservation

1. A grant or project Drive reading connection imports sources; Lab management → AI/Drive preservation configures the destination for copies and reports.
2. An administrator authorizes the preservation account and creates a dedicated folder or selects one accessible to the app.
3. Enable Markdown preservation, confirm data transmission and save. A folder URL alone does not prove authorization or upload success.
4. Inspect pending, successful and failed preservation counts and open actual files in Drive.

### 2. Preservation scope and recovery

1. Markdown is a readable body copy; research JSON and manifests preserve structured research information and versions.
2. This does not imply preservation of uncollected large files, video or model binaries. Operational DB dumps and authentication keys are separate.
3. For failures, ask the administrator to check account permission, folder access and storage.
4. Do not delete originals merely because a connection exists. Verify preservation and recovery before deciding retention policy.

### 3. Supply research intent for AI

1. Under Lab management → Research goals/AI context, record objectives, evaluation criteria, priorities and questions the report should answer.
2. Collect plans and proposals through project sources and mark relevant reference material.
3. Review report evidence and summaries and record where AI has inferred beyond the evidence.
4. AI consultation uses collected material you currently may access; check source connection and collection status first. Data storage is not model training, and human review remains necessary.

### Private reports and preserved copies

1. Only-me reports are excluded from new shared lab Drive preservation. Operational recovery backups, personal downloads and existing external copies are separate.

2. Moving content to Trash or disconnecting a source does not automatically delete published Notion/Drive originals.

3. Before migration or cleanup, verify actual copy locations and recovery with the administrator. Feature settings are not a server backup restoration procedure.

### Wiki originals and preservation status

1. Wiki bodies, revisions, proposals and original images use the lab Drive store. LUNA manages the index, structure and permissions.

2. Check Ready, Pending, Failed and last success in Backup. Completed analysis and a completed Drive copy are different states.

3. Reconnect the existing archive account for authorization errors. Changing folders requires planning access and migration of existing files; it is not a routine retry.

4. Personal AI memory, private conversations and Only-me reports have different preservation boundaries. Operational recovery requires verified DB, credentials and configuration backups; a help page cannot certify completion.

# 50. Troubleshooting and support

Troubleshoot access, authorization, collection and AI failures.

Open this guide online

## 1. Login, permissions and empty results

1. If returned to login, your session may have expired; sign in again.
2. For 403 or missing controls, check management permission. A 404 can also indicate deletion or unavailable access.
3. For empty lists, clear filters and verify approved, active participation.
4. For 502 or an unavailable site, record the URL and time and contact support by email or phone. Online support may also be unavailable during an outage.

## 2. Pending reports and collection errors

1. For collection errors, inspect the code and last successful time and follow the provider chapter.
2. For pending reports, check AI allowance, enablement and schedule status. Periodic source collection alone does not schedule reports.
3. Disabled schedules have no next run; completed reports remain. Avoid repeatedly resubmitting failures.
4. Include the error code, page URL, time and reproduction steps; never send tokens or passwords.

## 3. Submit a ticket and follow replies

1. Open Support immediately below Help in the profile menu. Contacts are public; online submission requires login.
2. Under New ticket enter a category, title, platform URL and details. Include steps, actual/expected results and the error code, then Submit ticket.
3. My tickets shows Received, Investigating and Resolved. New replies appear in the bell; open them for ticket details.
4. Add further details on the ticket page. A reply to a resolved ticket reopens it. Administrators answer and change status from All tickets. Email requires administrator Gmail setup and recipient opt-in. SMS and automatic error repair are not provided.

## 4. When another researcher cannot authorize Google

1. Read the error. If access is limited to approved test users, ask the connection administrator to check the intended Google email in the test-user list. Do not send passwords or tokens.

2. Administrators check audience, verification/testing state, test users and institutional policy. One successful account does not guarantee all accounts can authorize.
3. Confirm that the same Google account can directly open the original folder, file or calendar. Authorization and source access are separate.
4. After correcting policy, reconnect that card and verify collection. A `redirect_uri_mismatch` needs administrator configuration, not a different source folder URL.

Official OAuth guidance:

<https://developers.google.com/identity/protocols/oauth2/production-readiness/policy-compliance>

## 5. Slow pages and support resolution

1. Distinguish a loading section from an unreachable whole page. Discussions, profile work and summaries may load separately.
2. Narrow search and open only needed pages/threads. Use Previous/Next for favorites and work lists; distinguish No data from Loading.
3. Send Support the page URL, time, steps, expected result and actual result. Redact secrets and unrelated personal information in any screenshot.
4. New inquiries notify platform administrators; marking one resolved notifies its author. Open the inquiry from the alert to read the reply and status.
5. Repeat the affected steps to verify the fix. If unresolved, add the reproduction result to the existing inquiry.

## Reconnect an unavailable wiki store

1. If the list opens but the body says Wiki storage is unavailable, keep any unsaved text separately. The error alone does not mean the document was deleted.
2. Open Lab management → AI/Drive preservation → Connect Google Drive archive account. Sign in and consent with the account holding the existing wiki. This is separate from project Drive reading and personal file-picker authorization.
3. Do not create a new archive folder or change its address just to reconnect. Reopen the original page and check its body and revision history.
4. If it still fails, ask the administrator to check folder access with that account, preservation enablement, storage and revoked Google access. Do not overwrite an uncertain result with a blank page.
5. OAuth apps in Testing can have authorization expire after seven days. For recurring failures, ask the administrator to check publishing status and Google policy. Never include passwords or tokens in a ticket.

6. If consent returns “Another lab operation is in progress,” your existing connection was kept and the new authorization was not saved. Wait briefly, then start a new Google account connection from the same AI/Drive preservation settings. Do not refresh the failed callback or reuse its authorization URL. You do not need to stop ongoing research work.

Official guidance: <https://developers.google.com/identity/protocols/oauth2#expiration>

## When AI help cannot resolve an issue

1. Compare the cited guide with your screen. AI has not directly checked server state, permissions or connections.
2. If evidence is insufficient or the answer fails, open Review a Customer support draft. This did not resolve my issue opens the same draft after a completed answer.
3. Review and edit the question and AI answer, adding occurrence time, reproduction steps, and expected and actual results. Remove passwords, authentication codes and private research originals.
4. Choose Submit ticket to submit it yourself. Check Customer support for administrator replies and status. AI help does not submit, modify or mark issues resolved on your behalf.