The Thesis Pond

Onboarding Manual

Twelve interconnected tools for navigating your research journey — from proposal to defence.

OFFLINE.LTD · Version 2.0 · Screen Edition
Data stays on your device. Always.

Contents

Part I

Before You Begin

What the Pond Is (and Isn't)

You know that moment, three months into your research, when you realise your notes live in four different apps, your supervisor's feedback is buried in an email thread you can't find, and your literature review has become a folder of PDFs you half-remember reading? The Thesis Pond was built for that moment — and, ideally, to prevent it.

It is a single-page, offline-first web application that bundles twelve research-management tools into one coherent workspace. Half of them — the **Pond** side — handle the practical mechanics of a research degree: your brief, your literature, your writing progress, milestones, budget, supervisor meetings, and the decisions you made along the way. The other half — the **Forge** side — connect to local or cloud-based AI models to analyse your work through configurable research lenses, stress-test your arguments, and verify insights with dual-model synthesis.

What it is not: a word processor, a reference manager, or a substitute for your own critical thinking. Think of it more as a research cockpit — the instrument panel that helps you see where you are, what's overdue, and where the weak spots might be hiding.

> *The Pond works best when you visit it regularly — even for two minutes — rather than in marathon sessions before a deadline.*

How It Stores Your Data

Everything lives in your browser's localStorage. There is no server, no account, no cloud sync. When you close the tab, your data stays on that device, in that browser. This is a deliberate choice: your research is yours.

The trade-off is that you are responsible for backups. The Files tab lets you export your entire project as a JSON file at any time. Do this regularly. The app also keeps automatic version snapshots in localStorage, but these are not a substitute for a proper backup on a separate drive or cloud folder.

If you clear your browser data, your project goes with it. Treat the JSON export as your safety net.

System Requirements

The Thesis Pond runs in any modern browser — Chrome, Firefox, Safari, Edge — on desktop or mobile. No installation required; it's a single HTML file you can open from your hard drive, a USB stick, or a local server.

For the AI features (the Forge), you need either an Ollama instance running locally or an Anthropic API key. Without these, every Pond feature still works perfectly — the Forge tabs simply won't

produce analysis results.

Component	Requirement
Browser	Chrome 90+, Firefox 88+, Safari 14+, Edge 90+
Local LLM	Ollama running at localhost:11434 (optional)
Cloud LLM	Anthropic API key (optional)
Storage	~5 MB localStorage (grows with literature uploads)
Network	None — fully offline (except API calls to LLMs)

Your First Five Minutes

Open the HTML file. A welcome modal appears with two choices: **Load Demo Thesis** seeds the app with sample data so you can explore every feature immediately; **Start Fresh** gives you an empty workspace.

Whichever you choose, start with the **Research Brief** tab. Fill in your working title, degree type, discipline, institution, and dates. These details ripple through the entire application — the dashboard uses your start and submission dates to show a time-remaining indicator, and the analysis engine uses your research question and methodology to ground its feedback.

Next, hop to the **Literature Vault** and upload a few core papers. The app extracts text from PDFs automatically and sorts uploads into three weighted folders — Core, Supporting, and Working Notes — that feed the analysis engine.

Then wander. Click through the tabs. The Pond is designed for exploration, not instruction. You'll find your rhythm.

Understanding the Two Halves

The name gives it away: one *pond*, one *forge*. The metaphor is deliberate. A pond is still water — a place to gather things, let them settle, see your own reflection in them. The Pond side of the app is where you accumulate: documents, milestones, meetings, budgets, decisions. It doesn't judge. It holds.

A forge is heat and pressure — a place where raw material gets shaped into something stronger. The Forge side takes what you've gathered in the Pond and subjects it to analysis, critique, and structured questioning through AI lenses. It does judge, but constructively.

The two halves are connected by data: the Forge reads your Brief's research question, your Literature Vault's sources, and your selected lenses. The richer your Pond, the sharper your Forge. Neither half is useful without the other — a forge with no raw material produces nothing, and a pond with no heat stays tepid.

The Interface at a Glance

The app is organised as a tabbed interface. Across the top, you'll see the main tabs — Brief, Literature, Lenses, Analysis, Writing, Milestones, Supervisor, Progress, Budget, Dashboard, and Decisions — plus three utility tabs on the right: Settings (the gear icon), Files (the floppy disk), and Help (the question mark).

The header shows the app's name, the OFFLINE.LTD branding, and a theme toggle (the small circle button). If you have unsaved changes, a small dot appears next to the branding to remind you.

Each tab has a consistent layout: a title at the top, a brief description in italics, a horizontal divider, and then the tab's content — forms, lists, grids, or output areas. The design is intentionally understated. The focus is your research, not the interface.

Mobile & Progressive Web App

The Thesis Pond is responsive. On smaller screens, the tab bar scrolls horizontally, the two-column layouts stack vertically, and the progress wheel shrinks gracefully. Everything works on a phone, though the experience is naturally richer on a larger screen.

On supported devices and browsers, the app can be installed as a Progressive Web App (PWA). This means you can 'Add to Home Screen' on iOS or 'Install' on Android or Chrome, giving you an app icon, a splash screen, and a full-screen experience without the browser toolbar. The app manifest is embedded in the HTML — no additional files needed.

Even as a PWA, data stays in localStorage. The offline-first promise holds on every platform.

A Note on AI in Research

The Forge uses language models to analyse your research. This raises a question you should think about: where does AI assistance end and academic integrity begin?

The Thesis Pond takes a clear position: AI is a thinking partner, not a ghostwriter. The Forge never writes your thesis. It analyses, questions, and stress-tests — but the conclusions, the arguments, and the words are yours. Every output is clearly marked as AI-generated analysis, and nothing is auto-inserted into your writing.

Check your institution's AI policy. Most universities distinguish between AI as a drafting tool (often restricted) and AI as an analytical aid (usually permitted). The Thesis Pond sits firmly in the analytical aid category. But policies vary — and they are evolving fast.

> *Keep a record of how you use the Forge in your Decision Journal. If questioned, you will have a clear, timestamped trail of what AI contributed and what you decided.*

Part II

The Pond — Project Management

The Pond is the practical side of the Thesis Pond — nine tabs that handle the mechanics of keeping a research project alive, documented, and on schedule. None of these features require an AI connection. They work entirely offline, saving to your browser as you type.

Research Brief

The Brief is your project's identity card. It holds your working title, degree level, discipline, institution, key dates, research question(s), methodology, supervisor details, and contribution statement. Everything entered here is referenced by other tabs — the dashboard pulls your dates, the analysis engine reads your research question and methodology to contextualise its feedback.

You can change any field at any time. Research questions evolve, titles shift, methodologies pivot. The Brief is a living document, not a contract.

> *Write your contribution statement early, even if it's rough. Revisiting it every few weeks is one of the best ways to keep your research focused.*

The fields at a glance:

Field	Purpose	Used by
Working Title	Current thesis title	Exports, PDF header
Degree	PhD / Masters / Honours	Scope calibration in Forge
Discipline	Your field	Lens context
Institution	University name	Exports
Start / Submission / Defence	Key dates	Dashboard timeline, milestone warnings
Research Question(s)	Central + sub-questions	Analysis engine grounding
Methodology	Approach description	Arena stress-testing
Supervisor(s)	Names, roles, expertise	Supervisor log context
Contribution Statement	What this thesis adds	Contribution lens, Arena
Currency	€ / £ / $ / A$ / C$	Budget display

Literature Vault

The Vault is where your sources live. Drop in PDFs, plain text files, Markdown notes, or BibTeX entries. PDFs are auto-extracted using pdf.js in the browser — no server round-trip, no data leaving your machine.

Sources are sorted into three folders, each with a trust weight that influences how heavily the analysis engine leans on them:

Folder	Weight	What goes here
Core Sources	90%	Seminal papers, key theoretical works, primary data sources
Supporting Sources	60%	Secondary literature, methodological references, contextual reading
Working Notes	40%	Your own summaries, annotations, reading notes, draft fragments

BibTeX files get special treatment: each entry is parsed and turned into a reading-note template with fields pre-filled. This is especially useful if you manage your references in Zotero or Mendeley and want to bridge your bibliography into the Pond.

You can move sources between folders by dragging them, and delete individual sources with the × button. The 'Clear All Sources' button does exactly what it says — use it with intention.

 Upload your literature early and often. The more the Vault knows, the sharper the Forge's analysis becomes.

Walkthrough: Adding Your First Source

Click the dashed upload zone at the top of the Literature tab, or drag a file onto it. Suppose you drop in a PDF of a seminal paper — say, Smith and Jones (2021). The app will:

First, extract the text from the PDF using pdf.js, running entirely in your browser. Second, create a source entry in the Core Sources folder (the default destination). Third, display the source as a card showing the filename, file size, and extracted character count. The full text is now available to the analysis engine.

If you upload a .bib file instead, the app parses each BibTeX entry and creates a reading-note template for it — pre-filled with the author, title, year, and journal. This is a fast way to bootstrap your literature base if you already manage references in Zotero, Mendeley, or JabRef: export your library as .bib, drop it into the Vault, and you have a structured starting point.

To move a source between folders — say, from Core to Supporting — you can drag the source card. To remove a source, click the small × button on its card. Both actions take effect immediately.

How Trust Weights Shape Analysis

When the analysis engine builds its context for a query, it pulls text from the Vault. But not all sources are equal. A paper in Core Sources (90% weight) contributes proportionally more context than a working note (40% weight). This means the engine will lean more heavily on your primary literature and treat your own annotations as secondary inputs — which mirrors good academic practice.

You can exploit this. If you want the engine to pay special attention to a particular paper, put it in Core. If you want it to be aware of a reference but not dominated by it, put it in Supporting. If you want to inject your own interpretive frame into the analysis, write it up as a note and place it in Working Notes.

Writing Tracker

Add your thesis chapters or sections, each with a name and a word-count target. The tracker shows you at a glance how far along each chapter is — a set of cards with progress bars, plus aggregate statistics at the top: total words written, total target, and overall percentage.

You can update the current word count for any chapter as you go. The status field lets you mark chapters as 'Not started', 'Drafting', 'Revising', or 'Complete'. These statuses feed the Dashboard.

The Writing Tracker is deliberately simple. It's not trying to replace Scrivener or Word — it's the quick-glance overview you check before your supervisor asks 'so how's the writing going?'

Milestones

A flat list of deadlines with names, due dates, and completion status. Overdue milestones are flagged with a warning badge. Completed milestones stay visible (marked with a green tick) so you can see the arc of your progress. Summary statistics at the top show total, completed, pending, and overdue counts.

Examples of milestones you might track: ethics board approval, proposal defence, data collection complete, first draft of Chapter 3, pilot study submitted, external examiner confirmed.

> *Add milestones for soft deadlines too — 'finish reading Smith (2019)' or 'draft conference abstract'. Small wins compound.*

Working with Overdue Milestones

When a milestone passes its due date without being marked complete, it receives a warning badge. This is not a punishment — it is information. Research timelines slip. The important thing is knowing *how much* they have slipped and deciding consciously whether to extend, adjust scope, or push harder.

If you find yourself with multiple overdue milestones, resist the urge to delete them and pretend they never existed. Instead, update their due dates or mark them complete if they are no longer relevant, and log the reason in your Decision Journal. A pattern of overdue milestones is one of the earliest warning signs that your timeline needs recalibration — catching it here is far better than discovering it six months before submission.

Milestone Planning: A Practical Framework

When setting milestones, work backwards from your submission date. Start with the immovable deadlines — ethics approval, conference submission dates, your institution's hard deadlines. Then fill in the gaps with your own targets, allowing buffer time between each. A useful rule of thumb: whatever timeline you think a task will take, add 30%. Research consistently takes longer than expected, and building in buffer transforms 'always behind' into 'comfortably on track'.

Supervisor Log

Every meeting with your supervisor deserves a record. The log captures the date, what was discussed, what feedback was given, and what action items were agreed. Entries appear in reverse chronological

order.

This serves two purposes: it keeps you accountable to agreed actions, and it creates a paper trail. If there's ever ambiguity about what was said or decided — and in a multi-year project, there will be — the log is your reference.

The 'Export Supervisor Update' feature in the Files tab generates a clean summary you can send before or after meetings.

Getting the Most from Supervisor Meetings

The log works best as a three-part ritual. Before the meeting: review your previous log entry's action items and note which you have completed, which are in progress, and which slipped. During the meeting: take brief notes — even just keywords — on a piece of paper or your phone. After the meeting: transcribe your notes into the Supervisor Log within 24 hours, while the discussion is still fresh.

The action items field is especially important. Be specific: 'Revise methodology section to address external validity concern raised by Dr. Chen' is actionable; 'fix methodology' is not. Specific action items are easier to complete, easier to verify, and easier to discuss at the next meeting.

Over the course of a multi-year project, the supervisor log becomes a remarkably valuable document. It captures the evolution of your supervisor's thinking, the advice they gave at different stages, and the trajectory of your relationship. Some students include a version of this log in their thesis acknowledgements or methodology chapter.

Research Budget

Enter your total funding, then log expenses as they happen. Each expense has a category (Tuition/Fees, Equipment, Travel/Conferences, Software/Licences, Participant Costs, Books/Materials, Other), an item description, and an amount. The currency symbol is inherited from the Brief.

The summary cards at the top show total funding, total spent, remaining budget, and a colour-coded indicator: green when you're under 70% spent, amber between 70–90%, and red above 90%.

Decision Journal

Research is a long chain of decisions — narrowing scope, changing methodology, dropping a case study, switching theoretical frameworks. The Decision Journal logs each pivot with the date, what changed, the reasoning behind it, and the expected impact on timeline, methodology, or contribution.

This is more than record-keeping. Examiners love to see that you made deliberate, reasoned choices. A well-kept decision journal can strengthen your viva defence significantly.

> *Log decisions when they happen, not weeks later. Your future self will thank you for the context.*

What Makes a Good Decision Entry

A useful decision entry has three qualities: it is specific about what changed, honest about why, and realistic about the impact. Vague entries ('changed approach') are almost worthless when you return to them six months later. Compare these two entries:

Poor: 'Decided to change methodology.'

Good: 'Switched from structured interviews (n=30) to semi-structured interviews (n=15). Reasoning: pilot interviews revealed participants needed more space to elaborate on lived experience, and the structured format was constraining rich data. Impact: reduces sample size but deepens data quality; timeline shifts by 2 weeks; need to update ethics application.'

The second entry gives your future self — and your examiners — everything they need to understand the decision in context. It takes an extra two minutes to write. Those two minutes compound into hours saved during viva preparation.

Common Decision Types to Log

Type	Examples
Scope changes	Narrowing from 4 to 2 case studies; adding or dropping a research question
Method pivots	Switching analytical framework; changing sampling strategy; adding a pilot
Theoretical shifts	Adopting a new theoretical lens; abandoning a framework that did not fit
Timeline adjustments	Extending data collection; moving submission date; reordering chapters
Resource decisions	Purchasing software; hiring a transcription service; attending a conference
Supervisor-driven	Following supervisor advice to restructure; agreed change of direction

Progress Wheel

A visual self-assessment across eight dimensions: Literature, Method, Data, Writing, Wellbeing, Supervisor, Timeline, and Confidence. Each dimension is scored on a 1–10 slider, and the results are plotted as a radar chart — a spider web where you can immediately see which areas are strong and which need attention.

The centre of the wheel shows your overall average and a balance indicator. You can save snapshots at any time to track how your self-assessment changes over the life of your project.

Notice that Wellbeing is on the wheel alongside Literature and Method. That's intentional. A thesis is an endurance event, not a sprint. If your wellbeing score keeps drifting down, that's a signal worth paying attention to — no matter how the other numbers look.

Dimension	What it measures	Colour
Literature	Reading coverage, critical engagement with sources	Blue (#6a8aad)
Method	Methodological clarity, alignment with RQ	Green (#6aad6e)
Data	Data collection status, quality, access	Amber (#ad8a6a)
Writing	Drafts produced, revision stage	Purple (#8a6aad)
Wellbeing	Energy, motivation, mental health	Rose (#ad6a8a)
Supervisor	Relationship quality, meeting regularity	Teal (#6aada0)
Timeline	On schedule, ahead, behind	Olive (#adad6a)
Confidence	Self-belief in the project's value	Sage (#8aad6a)

Dashboard

The Dashboard is a read-only overview that pulls data from every other tab. At the top: summary cards showing time remaining until submission, overall writing progress, milestone completion rate, and budget status. Below that: writing progress per chapter, upcoming milestones, recent supervisor action items, and the latest analysis outputs.

You don't enter data here — you read it. Think of the Dashboard as the answer to 'where am I right now?' in a single glance.

Reading the Summary Cards

The top of the Dashboard shows four summary cards. Each one is a distilled metric drawn from a different part of the app:

Card	Shows	Source
Time Remaining	Days until submission, percentage of total time elapsed	Brief dates
Writing Progress	Total words vs. total target, overall percentage	Writing Tracker
Milestones	Completed / total count, overdue count if any	Milestones tab
Budget	Spent vs. total funding, colour-coded status	Budget tab

Below the cards, four sections drill deeper: a per-chapter writing progress breakdown with visual bars, a list of upcoming milestones sorted by due date, recent supervisor action items that may still need attention, and the last few analysis outputs from the Forge.

The Dashboard is the tab to open first thing in the morning and last thing before a meeting. It is the only tab that synthesises information from every other tab into a single view. If something is off — an overdue milestone, a budget running hot, a chapter falling behind — the Dashboard will surface it.

> *Pin the Dashboard as your default tab by opening it first whenever you launch the app. It takes three seconds to scan and instantly orients you.*

Part III

The Forge — AI-Powered Analysis

The Forge is where your research meets structured AI analysis. It does not write for you — it reads your Brief, your Literature Vault, and your chosen research lenses, then runs your questions through one or two language models to surface insights, contradictions, and blind spots.

All Forge features require a connected LLM backend — either Ollama running locally or an Anthropic API key entered in Settings. Without a backend, the Forge tabs are visible but non-functional.

Research Lenses & Presets

A lens is a perspective through which your research is evaluated. Each lens has a name, a core question, and a set of focus areas (weighted sub-topics). The analysis engine selects the active lenses to shape its prompts.

Three built-in preset packs give you ready-made lens configurations:

Preset	Lenses Included	Best For
Academic	Contribution, Rigour, Ethics, Scope, Impact	General research evaluation
Examiner	Argument, Literature, Method Fit, Presentation, Originality	Thesis defence preparation
Methodological	Design, Sampling, Validity, Reliability, Analysis	Method-heavy disciplines

You can toggle individual lenses on and off within a preset, adjust their weight sliders, or add entirely custom lenses with your own name, question, and focus areas. Custom lenses persist alongside the preset — they're yours to keep.

> *Create a custom lens for your specific examiner's known interests. If you know they care deeply about ecological validity, give the engine a lens for it.*

How focus areas work

Each lens contains 5–7 focus areas — clickable pills that narrow the analysis. When you click 'novelty' under the Contribution lens, the prompt sent to the LLM specifically targets how novel your work is, rather than evaluating contribution broadly. You can select multiple focus areas to combine perspectives.

Dual-LLM Synthesis

This is the Forge's signature feature. You write a research question or analysis prompt — 'Does my theoretical framework account for X?' or 'What tensions exist between source A and B?' — and two language models analyse it independently. Their outputs appear side by side in the dual-grid view.

Then a grounding step runs: the app sends both outputs back through the model with instructions to identify agreements, disagreements, and vault citations. The result appears in a highlighted synthesis box with a confidence meter.

Why two models? Because any single LLM can hallucinate, miss context, or lean into a particular interpretation. When two different architectures agree on a point, your confidence can rise. When they disagree, that's where the interesting questions live.

Choosing your models

The dropdown menus let you pick from locally available Ollama models (llama3.2, phi4, gemma2:9b) or Claude Sonnet via the Anthropic API. For dual verification, using two *different* models produces the most interesting results — for example, llama3.2 as Model A and phi4 as Model B. Running the same model twice still works but reduces the diversity of perspectives.

Single-Model Deep Dive

A simpler mode: one topic, one lens, one model. The Deep Dive is for when you want focused, extended analysis on a single aspect of your research. Enter a topic ('How robust is my sampling strategy?'), select a lens (e.g. Rigour), optionally click one or more focus areas (e.g. 'sample-design', 'validity'), and hit 'Go Deeper'.

The output is a single prose analysis rather than a dual comparison. Use this when you already know which angle you want to explore — the Dual Synthesis is for when you want to discover angles.

The Arena — Stress-Testing Your Thesis

The Arena is the Forge's most confrontational tool. It plays devil's advocate — deliberately attacking your reasoning, finding logical gaps, challenging your methodology, and questioning assumptions. Then, for every weakness it finds, it suggests how to defend against it.

You choose what to stress-test: the full thesis (research question + method + contribution), just the research question, just the methodology, just the contribution claim, or a custom passage you paste in.

Three intensity levels:

Intensity	Behaviour
Friendly Reviewer	Constructive, identifies gaps gently, assumes good faith
Rigorous Examiner	Thorough, expects strong evidence, asks hard questions
Hostile Opponent	Finds every weakness, no mercy, maximum pressure

> *Run the Arena at 'Hostile Opponent' intensity before your viva. If you can answer the Arena's attacks, you can answer most examiners.*

Walkthrough: A Typical Arena Session

Imagine you are writing a thesis on the impact of remote work on urban planning in mid-sized European cities. You have been using a mixed-methods design: surveys plus case studies. You have a nagging feeling that your contribution claim — that remote work fundamentally reshapes urban infrastructure demand — might be overstated. Time to enter the Arena.

Set the target to 'Contribution claim'. Set the intensity to 'Rigorous Examiner'. Select your preferred model and click 'Enter the Arena'. The output might look something like this (paraphrased):

"Your contribution claim assumes a causal link between remote work prevalence and infrastructure demand, but your methodology — surveys plus two case studies — cannot establish causation. You would need longitudinal data or a quasi-experimental design to support 'fundamentally reshapes'. Consider weakening the claim to 'is associated with shifts in' or strengthening the method to match the claim. Additionally, your case study selection (Amsterdam and Lisbon) introduces a Western European bias that limits generalisability..."

Each attack comes with a suggested defence. The Arena is not trying to discourage you — it is trying to make you stronger. The best response to an Arena session is not to feel defeated, but to feel prepared.

After a session, consider logging the key insights in your Decision Journal. If the Arena causes you to adjust your contribution claim, that is exactly the kind of methodological pivot examiners want to see documented.

Combining Forge Tools for Maximum Insight

The three Forge tools — Dual Synthesis, Deep Dive, and Arena — are most powerful when used together. A good sequence for exploring a complex question:

Start with a **Deep Dive** using a specific lens and focus area. This gives you a focused analysis of one angle. Then broaden with a **Dual-LLM Synthesis** — ask a wider question and let two models surface different perspectives. Finally, take your refined understanding into the **Arena** and stress-test it. The Arena will find gaps the other tools missed, because it is adversarial by design.

You can also reverse the sequence: start with the Arena to find your weakest points, then use Deep Dive to explore those weaknesses in detail, and finish with Dual Synthesis to see whether the fix you are considering holds up under cross-model scrutiny.

Working with Confidence Scores

After a Dual-LLM Synthesis, the confidence meter shows a percentage based on how closely the two models agreed. A high confidence score (green, 80%+) means both models reached similar conclusions. Medium (amber, 50–80%) means partial agreement with some divergence. Low (red, below 50%) means significant disagreement — which is exactly where you should dig deeper.

A low score is not a bad thing. It means the topic is genuinely complex or contested. That's valuable information for your research.

Understanding the Grounded Synthesis Output

The grounded synthesis — the highlighted output below the dual grid — is structured to help you act on what the models found. It typically contains three sections: agreements (where both models converged), disagreements (where they diverged), and vault citations (where the models referenced your uploaded literature). Pay special attention to the disagreements — they often reveal genuinely ambiguous aspects of your research that deserve more thought.

The vault citations are particularly useful because they ground the analysis in your actual sources. If the synthesis claims that 'Smith (2021) supports this interpretation', you can check whether Smith (2021) is actually in your Vault and whether the interpretation is fair. This is a built-in safeguard against hallucination — the models are prompted to cite from your literature, and you can verify those citations immediately.

Analysis History: Building a Record

Every Forge output — whether from Dual Synthesis, Deep Dive, or the Arena — is automatically saved to the Analysis History at the bottom of the Analysis tab. Each entry records the date, the prompt you used, the model(s) involved, and the full output text.

This history is more than a log. Over weeks and months, it becomes a chronicle of your analytical thinking — a record of what questions you asked, what the models found, and how your understanding evolved. When writing your methodology chapter, you can reference this history to document how AI-assisted analysis informed (but did not determine) your conclusions.

The analysis history is included in both the PDF and Markdown exports, giving you a portable record of your analytical journey.

Part IV

Files & Export

Saving and Loading State

The Files tab is your backup centre. **Export State (JSON)** downloads your entire project — every field, every source, every analysis result — as a single JSON file. **Import State** lets you load a previously exported file, restoring everything.

This is also how you move your project between devices or browsers. Export on your laptop, import on your office desktop. The JSON file is plain text and typically under 1 MB.

Exporting to PDF and Markdown

Two export formats are available for sharing or printing:

Export Full Project (PDF) generates a multi-page PDF containing your research brief, literature list, writing progress, milestones, supervisor log, decisions, budget summary, and analysis history. This is designed for handing to your supervisor or for your own records.

Export Full Report (Markdown) produces a Markdown document with the same content, suitable for pasting into Notion, Obsidian, or any other Markdown-aware tool.

Version History

The app automatically saves snapshots to localStorage at regular intervals. The Version History section in the Files tab shows these snapshots with timestamps. Click any snapshot to restore it. You can also create named versions manually — useful before major changes.

> *Name a version before every supervisor meeting. 'Pre-meeting 15 March' is far more useful than an auto-generated timestamp when you need to revert.*

Supervisor Sharing

The sharing feature generates a read-only link that contains your full project state encoded in the URL. Your supervisor can open it in any browser — no account needed, no server involved. The link is long (it literally contains your data), but it works.

This is a one-way snapshot, not a live link. Your supervisor sees the state at the moment the link was generated.

Supervisor Update Export

The 'Export Supervisor Update' button generates a concise summary document — a quick status report covering recent progress, upcoming milestones, open action items, and any decisions logged since the

last meeting. It is designed to be sent as an email attachment or printed before a meeting.

The update pulls data automatically from multiple tabs: writing progress from the Writing Tracker, open milestones from the Milestones tab, recent action items from the Supervisor Log, and recent decisions from the Decision Journal. You do not need to compile this information manually — the app assembles it for you.

> *Send the Supervisor Update 24 hours before your meeting. It gives your supervisor time to read it, which means your meeting can focus on discussion rather than status reporting.*

Walkthrough: A Complete Backup Routine

Here is a simple backup routine that takes under a minute and protects months of work:

First, open the Files tab and click 'Save Named Version'. Give it a meaningful name — 'Week 12 — lit review complete' is better than 'backup'. This creates a localStorage snapshot you can restore later without leaving the app.

Second, click 'Export State (JSON)'. Save the file to a cloud-synced folder — Google Drive, Dropbox, OneDrive, or even just email it to yourself. This is your off-device backup. If your hard drive fails or your browser data is wiped, this file is your lifeline.

Third, if you want a human-readable record, click 'Export Full Project (PDF)'. This is optional but useful for archival purposes — the PDF is a standalone document you can open anywhere.

Do this every Friday. Set a calendar reminder. The routine becomes automatic after two or three times, and the peace of mind is worth the sixty seconds.

Part V

Settings & Configuration

Connecting Ollama (Local LLMs)

Ollama is a tool that runs open-source language models on your own hardware. The Thesis Pond connects to it at http://localhost:11434 by default. If Ollama is running elsewhere on your network, change the URL in Settings.

To get started: install Ollama from ollama.ai, then pull a model — `ollama pull llama3.2` is a good starting point. Once it is running, the Thesis Pond's 'Test Connection' button will confirm the link.

Available local models include llama3.2, phi4, and gemma2:9b. You can pull any model Ollama supports — the app queries Ollama's API to discover what is available.

Step-by-Step: Installing Ollama

On macOS, download the installer from ollama.ai and drag it to your Applications folder. On Linux, run the one-line installer from the terminal: `curl -fsSL https://ollama.ai/install.sh | sh`. On Windows, download and run the .exe installer. Once installed, Ollama runs as a background service.

Pull your first model by opening a terminal and typing `ollama pull llama3.2`. This downloads the model weights (typically 2–8 GB depending on the model). You only need to do this once. To verify it is working, type `ollama list` — you should see your installed models.

Back in the Thesis Pond, go to Settings, ensure the provider is set to 'Ollama', and click 'Test Connection'. The status field should show 'Connected' along with a list of available models. If it fails, check that Ollama is running — you can verify by visiting http://localhost:11434 in your browser, which should display a simple confirmation message.

Choosing a Local Model

Model	Size	Strengths	Best For
llama3.2	~3 GB	Strong general reasoning, good at nuance	Default choice for most users
phi4	~2.5 GB	Fast, concise responses, good logic	Quick analyses, dual-model partner
gemma2:9b	~5 GB	Rich context handling, academic tone	Literature-heavy research

For dual-model verification, pairing two different architectures yields the most diverse insights. A good starting combination is llama3.2 as Model A and phi4 as Model B — they tend to approach problems differently, which is exactly what you want when looking for blind spots.

> *If your machine has limited RAM (8 GB or less), stick with phi4 — it is the lightest model and still produces useful analysis. If you have 16 GB or more, try gemma2:9b for richer outputs.*

Connecting the Anthropic API

If you prefer cloud-based models, select 'Anthropic' as the provider in Settings and enter your API key. This connects you to Claude Sonnet — Anthropic's capable, fast model. API usage is metered by Anthropic; the Thesis Pond does not add any cost on top.

Your API key is stored only in localStorage — it never leaves your browser except in direct API calls to Anthropic's servers. If you clear your browser data, you'll need to re-enter it.

Dual-Model Setup

The Dual-LLM Synthesis requires a second model. In Settings, the 'Second Model' section lets you configure a second provider — either a different Ollama model or Anthropic. Set it to 'Same as primary' if you want both analyses to use the same backend (though this reduces the diversity benefit of dual verification).

Theme & Display

The toggle in the header switches between dark and light themes. Your preference is saved. The dark theme is gentler on the eyes during late-night sessions. The light theme is optimised for daylight readability and printing.

On mobile, the interface is responsive — tabs scroll horizontally, and the layout adjusts for smaller screens. The app can also be installed as a Progressive Web App (PWA) on supported devices for an app-like experience without a browser chrome.

Part VI

Keyboard Shortcuts & Tips

Navigation Shortcuts

The Thesis Pond supports keyboard navigation for power users. All shortcuts use Ctrl (or Cmd on Mac) plus a number key:

Shortcut	Action
`Ctrl + 1`	Research Brief
`Ctrl + 2`	Literature Vault
`Ctrl + 3`	Research Lenses
`Ctrl + 4`	Analysis Engine
`Ctrl + 5`	Writing Tracker
`Ctrl + 6`	Milestones
`Ctrl + 7`	Supervisor Log
`Ctrl + 8`	Progress Wheel
`Ctrl + 9`	Budget
`Ctrl + 0`	Dashboard
`Ctrl + S`	Export State (JSON backup)
`Escape`	Close modal / welcome screen

Power-User Workflows

Here are three workflows that experienced users tend to settle into:

The Weekly Check-In

Every week, spend 10 minutes: update your Writing Tracker word counts, log any decisions made, check the Dashboard for overdue milestones, and save a Progress Wheel snapshot. This takes less time than it sounds, and the longitudinal data it creates is invaluable for your thesis introduction (where you often need to narrate your research journey).

Pre-Supervision Prep

Before meeting your supervisor: export a Supervisor Update, run the Arena at 'Rigorous Examiner' intensity on whatever you've been working on, and check your milestone status. Arrive with a printed update and a list of questions the Arena surfaced.

The Deep Work Session

When you're ready to think deeply about a specific aspect of your research: open the Lenses tab, select or create a custom lens that targets your concern. Go to Analysis, choose Deep Dive, and run 2–3 focused analyses from different focus areas. Then switch to Dual Synthesis with a broader question. Finish by running the Arena on the same topic. Save the JSON state when done.

The Literature Sprint

When you have a batch of new papers to process: upload them all to the Literature Vault in one go. Sort them into the appropriate folders. Then run a Dual-LLM Synthesis with a question like 'What are the key tensions between my existing Core sources and these new additions?' This surfaces contradictions and complementarities faster than reading each paper in isolation.

The Monthly Review

Once a month, take 30 minutes for a thorough review. Compare your current Progress Wheel snapshot to last month's. Read through your Decision Journal entries from the past four weeks. Check your budget trajectory. Run the Arena on your full thesis at 'Friendly Reviewer' intensity to get a gentle status check. This monthly ritual catches drift early — the kind of slow, invisible veering off course that only becomes obvious when it is too late to correct cheaply.

The Conference Prep

Presenting at a conference? Use the Arena at 'Hostile Opponent' on whatever you plan to present. The questions it generates are remarkably similar to the ones you will get from the audience. Run a Deep Dive on your weakest point. Export the results and bring them as notes — knowing the tough questions in advance transforms your confidence at the podium.

Troubleshooting

Problem	Solution
Forge says 'Backend not configured'	Go to Settings → select a provider → Test Connection
Ollama connection fails	Ensure Ollama is running: check with 'ollama list' in terminal
PDF upload doesn't extract text	pdf.js loads from CDN — check your internet connection once
Data disappeared	You may have cleared browser data. Import from your last JSON backup
Analysis is slow	Local models depend on your hardware. Try a smaller model like phi4
Confidence meter shows 0%	This usually means one model returned an error. Check the dual grid
App looks broken on mobile	Ensure you are not in extreme zoom. Try landscape orientation
BibTeX import creates blank entries	Check your .bib file for syntax errors — missing braces are common
Theme toggle does not persist	Cookies or localStorage may be blocked. Check browser privacy settings
Export PDF is empty	Fill in the Research Brief first — it anchors the export structure
Cannot install as PWA	PWA install requires serving over HTTPS or from localhost

Unsaved dot keeps appearing	This is normal — it shows whenever any field changes since last export

Part VII

Appendices

A — Lens Reference Table

All built-in lenses across the three preset packs, with their core questions and focus areas.

Lens	Preset	Core Question
Contribution	Academic	What new knowledge does this research add?
Rigour	Academic	Is the methodology sound? Would it survive peer review?
Ethics	Academic	Are ethical obligations met?
Scope	Academic	Is this finishable? Is the boundary clear?
Impact	Academic	Who benefits? What changes?
Argument	Examiner	Is the central argument coherent and well-supported?
Literature	Examiner	Does the review demonstrate mastery of the field?
Method Fit	Examiner	Does the method actually answer the research question?
Presentation	Examiner	Is the writing clear and at doctoral standard?
Originality	Examiner	Does this thesis show independent critical thinking?
Design	Methodological	Is the research design appropriate for the question?
Sampling	Methodological	Is the sample adequate and well-justified?
Validity	Methodological	Are the conclusions valid given the method?
Reliability	Methodological	Would another researcher get similar results?
Analysis	Methodological	Is the analytical approach rigorous and transparent?

B — Focus Areas by Lens (Academic Preset)

Lens	Focus Areas
Contribution	novelty, significance, originality, theoretical-advance, practical-value, field-positioning, knowledge-gap
Rigour	validity, reliability, sample-design, analysis-method, reproducibility, limitations-acknowledged, triangulation
Ethics	informed-consent, data-protection, vulnerable-populations, dual-use, positionality, cultural-sensitivity, institutional-c
Scope	research-boundary, time-feasibility, resource-constraints, depth-vs-breadth, inclusion-criteria, exclusion-criteria, com
Impact	academic-impact, societal-impact, policy-relevance, practitioner-value, future-research, dissemination-strategy, citatio

Focus Areas by Lens (Examiner Preset)

Lens	Focus Areas
Argument	thesis-statement, logical-structure, evidence-quality, counter-arguments, consistency, conclusions-follow, reader-pers
Literature	coverage, critical-analysis, synthesis, gap-identification, theoretical-framework, currency, seminal-works
Method Fit	alignment, justification, alternatives-considered, data-adequacy, analytical-framework, reflexivity, limitations
Presentation	clarity, structure, academic-voice, referencing, figures-tables, proofreading, formatting
Originality	independent-thinking, creative-synthesis, new-perspective, challenging-assumptions, novel-connections

Focus Areas by Lens (Methodological Preset)

Lens	Focus Areas
Design	appropriateness, alignment-with-rq, philosophical-underpinning, design-flexibility, scalability, replication-potential
Sampling	sample-size, selection-criteria, representativeness, access-feasibility, bias-mitigation, saturation
Validity	internal-validity, external-validity, construct-validity, ecological-validity, face-validity, member-checking
Reliability	inter-rater, test-retest, consistency, audit-trail, dependability, confirmability
Analysis	coding-framework, thematic-saturation, statistical-power, software-tools, transparency, reflexive-practice

C — Budget Categories

Category	Typical Items
Tuition / Fees	Annual fees, bench fees, lab access charges
Equipment	Laptop, recording devices, lab equipment, consumables
Travel / Conferences	Conference registration, flights, accommodation, poster printing
Software / Licences	NVivo, SPSS, Adobe Creative Suite, journal access
Participant Costs	Incentives, gift cards, transcription services
Books / Materials	Physical copies, special collections access, interlibrary loans
Other	Printing, binding, professional development courses

D — Glossary

Arena — The stress-testing tool in the Forge that plays devil's advocate against your research arguments.

Brief — The Research Brief tab — your project's core identity: title, RQ, method, contribution.

Confidence Meter — A percentage shown after Dual-LLM Synthesis indicating how closely two models agreed.

Deep Dive — Single-model, single-lens focused analysis in the Forge.

Dual-LLM Synthesis — Running two language models in parallel on the same prompt, then synthesising agreements and disagreements.

Focus Area — A sub-topic within a research lens that narrows the analysis (e.g. 'novelty' within Contribution).

Forge — The AI-powered analysis side of the Thesis Pond — Lenses, Analysis, and Arena.

Grounded Synthesis — The final output of Dual-LLM Synthesis, combining both model outputs with vault citations.

Lens — A configurable research perspective (e.g. Rigour, Ethics, Scope) used to evaluate your work.

localStorage — The browser storage mechanism where all Thesis Pond data is kept.

Ollama — An open-source tool for running LLMs locally on your own computer.

Pond — The project-management side of the Thesis Pond — Brief, Literature, Writing, Milestones, etc.

Preset — A pre-configured pack of research lenses (Academic, Examiner, or Methodological).

Progress Wheel — A radar chart showing your self-assessment across eight research dimensions.

Trust Weight — The percentage weight (40–90%) given to sources in different Literature Vault folders.

Vault — The Literature Vault — where uploaded sources are stored and organised.

E — Data Architecture Quick Reference

Understanding how data flows through the Thesis Pond helps you use it more effectively. Below is a simplified view of how the app's components connect.

Data Source	Feeds Into	How
Brief → RQ, Method	Analysis Engine, Arena	Included in LLM prompts as context
Brief → Dates	Dashboard	Time-remaining calculation
Brief → Currency	Budget	Currency symbol display
Literature Vault	Analysis Engine	Source text weighted by folder trust level
Lenses	Analysis Engine, Deep Dive	Shapes prompt questions and focus areas
Writing Tracker	Dashboard	Aggregate word counts and chapter status
Milestones	Dashboard	Overdue detection and upcoming list
Supervisor Log	Dashboard, Export	Recent actions shown on Dashboard
Progress Wheel	Snapshots	Saved as timestamped records
Budget	Dashboard	Spending vs. funding summary
Decisions	Export	Included in full project PDF/MD export
Analysis History	Dashboard, Export	Recent results and full log in exports

F — Security & Privacy Notes

The Thesis Pond stores all data in your browser's localStorage. No data is sent to any server unless you explicitly use the Anthropic API for analysis — in which case, only the specific prompt (including your research question and relevant source excerpts) is transmitted to Anthropic's servers.

When using Ollama, all processing happens on your local machine. Your data never leaves your hardware.

The shared supervisor link encodes your full project state in the URL itself — it is not stored on any server. However, anyone with the link can view the data, so share it consciously.

Your Anthropic API key, if entered, is stored only in localStorage. It is transmitted directly to Anthropic's API endpoint and nowhere else.

G — Frequently Asked Questions

Can I use this for my undergraduate dissertation?

Absolutely. Although the Thesis Pond was designed with doctoral and masters students in mind, undergraduates working on a significant dissertation will find every feature useful. The degree selector in the Brief includes an 'Honours / 4th year' option, which calibrates the analysis engine's expectations accordingly.

Can multiple people use the same project?

Not simultaneously. The Thesis Pond is single-user by design — there is no server to coordinate multiple sessions. However, you can export your project as JSON, send it to a collaborator, and they can import it on their machine. Changes made on separate machines will need to be reconciled manually.

How much data can the Vault hold?

localStorage typically allows 5–10 MB per origin (domain), depending on the browser. In practice, this is enough for 20–40 substantial PDFs (text-extracted, not the raw binary). If you hit the limit, consider removing older working notes or sources you no longer need. The Files tab shows your current storage usage.

Will my data survive a browser update?

Yes, in virtually all cases. Browser updates do not clear localStorage. The only scenarios that delete your data are: manually clearing browser data (including site data), uninstalling the browser, or using a privacy-focused browser mode that wipes storage on close. As always, regular JSON exports are your best insurance.

Can I use the Forge without any AI backend?

The Forge tabs will be visible but non-functional — clicking 'Analyse' or 'Go Deeper' will not produce results. Every Pond feature (Brief, Literature, Writing, Milestones, Supervisor, Budget, Decisions, Progress, Dashboard) works perfectly without any AI connection.

Does the analysis engine remember previous analyses?

Each analysis run is independent — the engine does not carry memory between runs. However, all results are saved in the Analysis History, which you can review at any time. The context for each run comes from your Brief, Literature Vault, and selected lenses — not from previous analyses.

Can I customise the eight dimensions in the Progress Wheel?

Not in the current version. The eight dimensions — Literature, Method, Data, Writing, Wellbeing, Supervisor, Timeline, and Confidence — are fixed. They were chosen to cover the full spectrum of a research degree, including the often-neglected human dimensions. Custom dimensions may be added in a future version.

What happens if I switch from Ollama to Anthropic mid-project?

Nothing breaks. Your project data is independent of the AI backend. Analysis history from previous runs (with either backend) is preserved. The only difference is which model processes your next query. You can switch back and forth freely.

Can I self-host the Thesis Pond on a local server?

Yes. Since it is a single HTML file, you can serve it from any local HTTP server — Python's `python -m http.server`, Node's `npx serve`, or any lightweight web server. This can be useful in a lab or research group where multiple people want access. Each person's data will still be stored in their own browser's localStorage.

Is there a mobile app?

Not a native app, but the PWA installation gives you an app-like experience on both iOS and Android. Add to Home Screen from Safari (iOS) or Chrome (Android), and you will get an icon, splash screen, and full-screen interface. Data syncing between devices still requires manual JSON import/export.

H — Mapping the Pond to Research Phases

Different stages of a research degree demand different tools. Here is a guide to which Thesis Pond features are most useful at each phase.

Phase	Primary Tabs	Secondary Tabs
Proposal / Design	Brief, Lenses, Literature	Milestones, Decisions
Ethics / Approval	Milestones, Supervisor	Brief, Decisions
Literature Review	Literature, Analysis (Deep Dive)	Writing, Lenses
Data Collection	Milestones, Progress, Budget	Supervisor, Decisions
Analysis / Coding	Analysis (Dual + Deep Dive)	Writing, Literature
Writing Up	Writing, Dashboard	Supervisor, Decisions, Budget
Revision / Editing	Writing, Arena	Supervisor, Literature
Pre-Viva Prep	Arena, Analysis, Dashboard	Progress, Decisions
Post-Viva Corrections	Writing, Supervisor, Milestones	Decisions, Dashboard

Most researchers will not use every tab with equal intensity at every phase. The table above is a starting point — your own workflow will evolve as your project does. The Dashboard is useful at every phase, which is why it sits at the centre of the app.

I — Designing Your Own Lenses

The built-in presets cover common research evaluation angles, but your specific project may benefit from custom lenses that target your unique concerns. Here is a framework for designing effective custom lenses.

Step 1: Identify the concern

What keeps you up at night about your thesis? Is it whether your theoretical framework holds together? Whether your data is sufficient? Whether your contribution is genuinely novel? Name the concern in plain language.

Step 2: Frame it as a question

Turn that concern into a single, evaluative question. Good lens questions are open-ended, evaluative (not yes/no), and focused on one dimension. Examples: 'Is the theoretical framework internally consistent?', 'Does the sampling strategy introduce systematic bias?', 'Would a practitioner find actionable insights in this research?'

Step 3: Define focus areas

Break the question into 4–7 sub-aspects. These become the clickable focus pills in the Analysis tab. Use hyphenated compound terms — 'framework-coherence', 'assumption-transparency',

'cross-case-consistency'. Each focus area should be narrow enough to generate a targeted analysis when selected.

Step 4: Test and refine

Add the lens, run a Deep Dive with it, and read the output critically. Does the analysis address your actual concern? If the output feels too generic, tighten the core question. If it misses important angles, add more focus areas. Custom lenses are meant to evolve with your project.

> *Create a lens named after your most critical examiner — frame the question as what they would ask. It is remarkably effective at surfacing blind spots.*

Example Custom Lenses

Name	Core Question	Focus Areas
Positionality	How does my background shape this research?	reflexivity, bias-awareness, insider-outsider, privilege-dynamics, cultura
Scalability	Could this research be replicated at larger scale?	resource-requirements, generalisability, automation-potential, data-acce
Practitioner Value	Would a professional in this field find this useful?	actionability, implementation-ease, cost-benefit, stakeholder-relevance,
Supervisor Lens	What would Prof. Martinez ask about this?	theoretical-depth, methodological-rigour, writing-clarity, argument-flo
Counter-Narrative	What voice is missing from this research?	excluded-perspectives, power-dynamics, representation, alternative-fra

These are starting points, not prescriptions. The best custom lenses are the ones that make you slightly uncomfortable — they ask the question you have been avoiding. That discomfort is precisely where growth happens, and where examiners will inevitably probe.

A well-designed set of custom lenses — three or four, tailored to your project — can transform the Forge from a generic analysis tool into something that feels like a conversation with a knowledgeable colleague who knows your work intimately.

J — Degree-Specific Guidance

The Thesis Pond adapts to different degree levels through the Brief's degree selector, but the way you use the app should also vary based on the scale and expectations of your project.

PhD / Doctorate

You are in this for three to five years. The Pond becomes your institutional memory — the single place where the full arc of your research is documented. Prioritise the Decision Journal (examiners at this level expect to see rigorous, documented reasoning), the Literature Vault (build it deep, not just wide), and the Arena (stress-test your contribution claim relentlessly). Use all three Forge tools regularly. Save Progress Wheel snapshots monthly to capture your development narrative.

The Dual-LLM Synthesis is particularly valuable for doctoral work because it surfaces the kind of nuanced disagreements that live at the frontier of knowledge — exactly where your thesis needs to sit.

Masters (Thesis)

You have 12–18 months and a shorter document to produce. Focus on the Brief (get your RQ tight early), the Writing Tracker (chapter targets keep you honest), and Milestones (you have less margin for slippage). The Arena at 'Friendly Reviewer' intensity is a good fit — rigorous enough to improve your work, not so intense that it overwhelms a shorter project. The Examiner preset pack is your best starting point for lenses, since masters evaluation criteria often mirror those lenses closely.

Masters (Dissertation)

Often a shorter, more focused project than a full thesis — sometimes 15,000 to 20,000 words. The Writing Tracker and Milestones are your core tools. Use the Forge selectively: one or two Deep Dives on your key concern, and a single Arena session before submission. The Budget tab may be less relevant unless your programme requires funded research.

Honours / 4th Year

This is often your first major independent research project. The Brief is essential — writing out your research question, methodology, and contribution forces clarity that the project desperately needs at this stage. The Supervisor Log is especially valuable because the supervisor relationship is new to you and documenting it helps you learn how to be supervised. Use the Forge gently: the Deep Dive on single topics will be most useful, and the Arena at 'Friendly Reviewer' can help you anticipate examiner questions without being overwhelming.

> *Whatever your degree level, start using the Thesis Pond in your first month — not your last. The app's value compounds with time. A project documented from day one tells a richer story than one reconstructed from memory the week before submission.*

K — Colophon

The Thesis Pond was designed and built by OFFLINE.LTD. This manual was typeset in Crimson Text, an open-source serif typeface designed by Sebastian Kosch, chosen for its warmth, readability, and quiet authority — the same qualities we aspire to in research writing.

The colour palette centres on a muted teal (#2E7D6F) that recalls still water — steady, calm, deep enough to hold complexity. The dark theme uses warm off-whites on near-black backgrounds to reduce eye strain during late-night writing sessions. The light theme reverses this with parchment-toned backgrounds for comfortable daylight reading and clean printing.

The application itself is a single HTML file — no build tools, no dependencies, no server. This is a philosophical choice as much as a technical one. Research tools should be durable. They should work in five years without maintenance. A single file that runs in any browser comes as close to that ideal as current technology allows.

If you are reading this manual, you are probably somewhere in the middle of a research journey. It is hard work. There will be weeks when the data doesn't cooperate, when the writing feels impossible, when the gap between what you know and what you need to know seems uncrossable. That is normal. That discomfort is the shape of learning.

The Thesis Pond cannot do the hard work for you. But it can help you see where you are, remember where you have been, and think more clearly about where to go next. That is enough. That is the job.

The Thesis Pond was built by researchers, for researchers.

Twelve tools. One file. Your data, on your device, under your control.

Now close this manual and go write something.

OFFLINE.LTD · v2.0
thesispond.offline.ltd