
ResearchForge

Onboarding Manual

*A thinking tool for researchers who believe the best questions
are the ones that survive the hardest scrutiny.*

Version 1.0 · 2026

Screen Edition — optimised for dark displays

*“The purpose of science is not to prove yourself right.
It is to become less wrong, one careful question at a time.”*

Contents

I What Is This Place?

The shape of ResearchForge

What it is not

Who it is for

II First Light

Opening the interface

The sidebar and the main stage

Choosing a backend

III The Source Vault

Uploading your materials

Three folders, three trust tiers

Chunking, indexing, and the weight sliders

IV Analytical Lenses

Presets: five ways of seeing

Custom lenses

Weight and activation

V The Synthesis Engine

Writing a research question

Reading the output

Contradictions, citation chains, and gaps

VI Citation Map

Tracing a claim through your sources

The citation graph

Quick Scan and the contradiction detector

VII Deep Dive

Choosing a lens and a focus

The specialist reviewer

VIII Storm

- Adversarial pressure-testing
- Three passes of increasing intensity
- Reading the verdict

IX Notes, History & Export

- Keeping a research diary
- Export and import state
- Markdown exports

X Settings & Backends

- Ollama, WebAI, StudioLM, Anthropic
- The fallback chain
- Vignette and visual tuning

XI Workflows & Recipes

- Literature review workflow
- Hypothesis pressure-testing
- Grant preparation

XII Philosophy of the Forge

- Why pressure-testing matters
- The limits of AI-assisted research
- Epistemic humility as a practice

CHAPTER I

What Is This Place?

You have a stack of papers. Some agree with each other. Some quietly do not. This is the tool that helps you find out which is which.

Imagine you are alone in a library at 2 a.m., surrounded by printed papers. Some of them contain findings that contradict the paper right next to them, and neither author seems to know. You have a highlighter, a coffee, and a gnawing suspicion that something in your literature review does not hold together.

ResearchForge is the digital version of that 2 a.m. session — except it reads faster than you do, and it does not get tired. It is a synthesis and pressure-testing studio that takes your research sources, examines them through configurable analytical lenses, and flags where claims support each other, where they collide, and where the evidence quietly vanishes.

The Shape of the Forge

At its core, ResearchForge is nine tabs and a sidebar. The tabs are stages in a workflow: you upload sources, choose how to interrogate them, run synthesis, trace citations, dive deeper on specific questions, and finally storm-test your conclusions against adversarial pressure. The sidebar keeps your research threads — think of them as parallel lines of inquiry you can switch between without losing your place.

Every analysis is powered by a language model running either locally on your machine or via the Anthropic API. The important thing: **ResearchForge is a thinking tool, not a peer reviewer.** Its outputs are suggestions, provocations, and flags — never conclusions. You remain the scholar. The forge just heats the metal; you do the shaping.

A quiet reminder. Every AI output in ResearchForge is suggestive, never authoritative. Claims should always be traced back to primary sources. Apply your own scholarly judgment to every output — the tool is only as trustworthy as the questions you ask it.

What It Is Not

It is not a citation manager — it does not format your references. It is not a writing assistant — it does not draft your paper for you. And it is not a replacement for actually reading the sources yourself. Think of it more like a very patient colleague who has read everything in your stack and is willing to play devil's advocate at any hour.

Who It Is For

Graduate students wading through their first proper literature review. Post-docs pressure-testing a hypothesis before committing to a year of lab work. Grant writers checking that their innovation claim actually holds up. Anyone who has ever wondered, "*Wait, does this paper actually contradict that one?*" and wished for a faster way to find out.

If you have never asked that question, this tool will make you start. That is rather the point.

CHAPTER II

First Light

Opening the forge for the first time. The glow of a dark screen, the quiet hum of possibility.

Open the **Research_Forge.html** file in any modern browser. That is it. No installation, no build step, no dependency chain. A single file that contains the entire studio.

The first thing you will see is a dark interface — teal accents against near-black — with a sidebar on the left and a tabbed workspace on the right. There is a gentle noise texture across the screen and a vignette darkening the edges. This is deliberate: the interface is designed for long reading sessions, the kind where you forget what time it is.

The Sidebar

The sidebar holds your **research threads**. Each thread is a named line of inquiry — "catalyst temperature effects," "replication concerns," "grant draft 3." You can create new threads, search through existing ones, and switch between them without losing any analysis. Think of threads as parallel notebooks.

The Top Bar

Along the top you will find Export (saves your entire state as JSON), Import (restores it), a theme toggle (the half-moon icon switches between dark and light mode), and a backend status indicator. That little dot on the right tells you whether your AI backend is connected. Green means go. Yellow means it is checking. Red means something needs attention.

Choosing a Backend

ResearchForge supports four AI backends, and you do not need to choose just one. The system has a fallback chain — if your primary backend goes down mid-analysis, it silently switches to the next configured one.

Backend	Type	Best For	Port
---------	------	----------	------

Ollama	Local	Privacy, offline work, free	11434
WebAI	Local	OpenAI-compatible servers	8080
StudioLM	Local	LM Studio users	1234
Anthropic	Cloud	Highest quality synthesis	—

No backend? No problem. Without any AI backend, you can still manage sources, run the client-side Quick Scan contradiction detector, use citation mapping, keep notes, and export your work. The AI features light up when a backend connects — everything else works regardless.

CHAPTER III

The Source Vault

Papers, datasets, lab notes, competing theories. All of them go here, weighted by how much you trust them.

The Source Vault is where you feed the forge. Drag and drop .txt, .md, or .pdf files into the upload zone, or click to browse. ResearchForge reads them, splits them into chunks, and sorts them into one of three trust-tier folders based on filename heuristics.

Three Folders, Three Trust Tiers

Every source lands in one of three folders. Each folder has a **trust weight slider** — a percentage that tells the synthesis engine how much epistemic confidence to place in materials from that tier. You can adjust these at any time.

Folder	Default Weight	What Goes Here
Papers	80%	Peer-reviewed articles, published studies, preprints, meta-analyses
Methods	60%	Protocols, datasets, supplementary materials, replication notes
Working	40%	Lab notes, drafts, preliminary data, personal observations

The folder assignment is based on keywords in the filename — words like "paper," "journal," or "review" route to Papers; "protocol," "dataset," or "analysis" route to Methods; everything else goes to Working. You can adjust manually by re-uploading with a different name.

Chunking and Indexing

When you upload a file, ResearchForge splits it into chunks of roughly 800 characters with 200-character overlap. This overlap ensures that ideas spanning paragraph boundaries are not lost. Think of each chunk as a quote card in a physical research file.

Clicking **Index All Sources** sends each chunk to your embedding model, generating vector representations. These vectors are what enable the RAG (Retrieval-Augmented Generation)

system to find relevant context when you run a synthesis. Without indexing, the AI still works — it just cannot search your sources intelligently. With indexing, it reads through your materials like someone who has actually studied them.

Indexing requires a local backend. Embeddings are generated by models like nomic-embed-text running on Ollama, WebAI, or StudioLM. The Anthropic API does not provide embeddings — if you use Claude as your primary backend, keep a local server running alongside it for indexing.

Source Statistics

The stats bar at the top of the Source Vault shows you four numbers: total sources, total chunks, percentage indexed, and number of papers. These update in real time as you add, remove, or index materials. A fully indexed vault is a responsive vault — synthesis runs faster and retrieves more relevant context.

CHAPTER IV

Analytical Lenses

Five presets, infinite custom lenses. Each one is a different pair of glasses through which the forge examines your research.

A lens is a mode of interrogation — a specific question or perspective that the synthesis engine applies to your research materials. If the Source Vault is *what* you are studying, lenses are *how* you study it.

The Five Presets

Epistemic Rigor

Methodology, evidence quality, statistical validity, causal inference, falsifiability. This is the preset for researchers who want to know: *"Is this actually true, and how would we know?"*

Reproducibility

Replication, transparency, materials, data access, protocol detail. The preset that asks: *"Could someone else do this again and get the same result?"*

Interdisciplinary

Cross-field connections, translation, boundary concepts, methodological import. For those moments when you suspect that a biologist and a physicist are describing the same phenomenon in different languages.

Funding Feasibility

Grant readiness, significance, innovation, approach, investigator strength. The hard-nosed preset that asks what a review panel would ask.

Meta-Analysis

Synthesis quality, heterogeneity, publication bias, effect estimation, GRADE. The preset for researchers evaluating bodies of evidence rather than individual studies.

Each preset loads five specialised lenses, each with its own core question and a set of focus sub-topics. You can activate a preset with one click, then fine-tune individual lens weights or toggle specific lenses off.

Custom Lenses

If the presets do not cover your angle, build your own. A custom lens needs just two things: a name and a core question. The core question is the lens's soul — it defines what the AI looks for when reading through your sources. Make it specific. *"Is this good?"* will get you vague answers. *"What unstated assumptions does this model make about the distribution of error terms?"* will get you useful ones.

Weight and Activation

Every lens has two controls: a **weight slider** (0–100%) and an **on/off toggle**. The weight determines how much emphasis that lens receives in synthesis outputs — a lens at 100% dominates the analysis, while one at 20% contributes a quieter voice. The toggle lets you deactivate a lens without losing its configuration.

In practice, you will find yourself adjusting weights as your understanding of the material evolves. Early in a review, you might weight Methodology heavily. Later, as you move toward writing, Funding Feasibility might climb. The lenses shift with your thinking — that is by design.

CHAPTER V

The Synthesis Engine

The heart of the forge. Where your research question meets your sources and the lenses do their work.

The Synthesis tab is where everything converges. You type a research question — or a hypothesis, or a prompt for literature synthesis — and the engine reads your indexed sources through every active lens, then produces a structured analysis.

Writing a Good Research Question

The quality of the output is proportional to the specificity of the input. Here are some patterns that work well:

Evidence synthesis: "What is the current evidence for X? Where do the sources agree and disagree?"

Contradiction hunt: "How do these three models of Y contradict each other?"

Gap analysis: "What methodological gaps exist in the literature on Z?"

Pressure test: "Pressure-test this hypothesis against the uploaded papers."

Reading the Output

The synthesis output is structured in sections. First, you get a **per-lens analysis** — one focused paragraph for each active lens, applying its specific perspective to your question. Then come the cross-cutting sections:

Contradictions flag where sources or claims directly conflict, naming specific claims, the sources that make them, and what evidence or experiment would resolve the disagreement.

Citation Chains identify which claims are upstream (foundational) and which are downstream (derived), showing where a chain of reasoning is only as strong as its weakest link.

Gaps catalogue what the literature has not studied — unanswered questions, under-represented populations, conditions that have not been tested.

Finally, the engine highlights the most **surprising finding** or the claim most vulnerable to pressure-testing.

Evidence confidence markers. Look for [STRONG EVIDENCE], [WEAK EVIDENCE], and [CONTRADICTION] tags in the output. These are the engine's way of flagging its own confidence levels. Treat them as starting points for your own investigation, not as verdicts.

CHAPTER VI

Citation Map

Trace a single claim through your entire source collection. See who supports it, who contradicts it, and who quietly changes the subject.

The Citation Map is for precision work. Instead of asking a broad research question, you enter a specific claim — "Higher reaction temperatures improve catalyst yield" — and the engine traces how each of your sources relates to that claim.

Four Relationships

Every source is mapped to one of four relationship types:

Type	Meaning
Supports	The source provides evidence in favour of the claim
Contradicts	The source provides evidence against the claim
Extends	The source builds on the claim with additional findings
Qualifies	The source adds boundary conditions or limitations

The Citation Graph

After a trace completes, ResearchForge renders a force-directed graph (using D3.js if available, or a static radial layout as fallback). Your claim sits at the centre, and each source orbits around it. The colour of each connection indicates the relationship type, and the size of each node reflects the evidence strength — strong, moderate, or weak.

Hover over any node to see the key evidence extracted from that source. It is a map in the truest sense — a way of seeing the landscape of evidence around a single idea.

Quick Scan

CHAPTER VII

Deep Dive

Choose a single lens. Choose a single focus within that lens. Then go deeper than the synthesis engine ever could.

Where the Synthesis Engine gives you breadth — every active lens applied to a question — the Deep Dive gives you depth. You select one lens, one focus sub-topic within that lens, and enter a specific query. The AI then operates as a **specialist reviewer**, applying only that particular perspective with full intensity.

For example: you might select the Epistemic Rigor preset's Evidence Quality lens, choose the "power-analysis" focus, and ask: "Is the reported effect size in Chen 2024 large enough to detect with their sample size?" The Deep Dive would then produce a focused, expert-level analysis of that specific question.

Focus Sub-Topics

Each lens comes with a set of focus pills — predefined sub-topics that narrow the analysis. The Methodology lens, for instance, offers: study-design, control-conditions, randomisation, blinding, sample-selection, measurement-validity, and operationalisation. Clicking a pill tells the AI to concentrate on that specific aspect.

These are not decorative — they meaningfully change the output. A Deep Dive on "methodology" with focus "randomisation" will produce a very different analysis than the same lens focused on "measurement-validity."

When to use Deep Dive vs. Synthesis. Use Synthesis for broad landscape views across all your lenses. Use Deep Dive when you have found a specific crack in the wall and want to pry it open further. The two complement each other — Synthesis finds the interesting questions, Deep Dive answers them.

CHAPTER VIII

Storm

Adversarial pressure-testing. Three passes, escalating intensity. Your conclusions either survive the storm or they crack under it.

The Storm is the forge's most aggressive tool. You paste in a conclusion, a research summary, or a synthesis output, select an intensity level, and the engine runs three sequential passes of adversarial pressure-testing — each one designed to find weaknesses the previous pass missed.

Three Passes

Pass 1 — Evidence Audit. The engine examines every factual claim in your text, checking whether it is supported by your sources, contradicted by them, or unsubstantiated. It flags weak evidence, missing citations, and overgeneralisations.

Pass 2 — Logic and Methodology. The engine attacks the reasoning structure. It looks for logical fallacies, unsupported causal claims, selection bias in your choice of sources, and assumptions that have not been made explicit. It thinks like a hostile reviewer who is technically skilled and slightly annoyed.

Pass 3 — Steelman Counter-Argument. The engine constructs the strongest possible argument *against* your conclusions, using evidence from your own sources. This is not straw-manning — it is the genuine best case that a well-informed opponent could make. If your argument survives this, it is robust.

Intensity Levels

Level	Behaviour
Standard	Balanced analysis, fair scrutiny, constructive tone
Aggressive	Actively seeks flaws, challenges every claim, adversarial
Hostile	Full adversarial mode — assumes bad faith, attacks weakest links ruthlessly

The Verdict

After all three passes complete, the Storm produces a verdict — a summary assessment of how well your conclusions held up. Three possible outcomes:

Survived the Storm. No critical vulnerabilities found. This does not mean your conclusions are correct — it means the obvious failure modes have been checked.

Partially Survived. Some conclusions held, but at least one critical weakness was found. The analysis needs reinforcement before you build further.

Cracked Under Pressure. Multiple critical vulnerabilities. The load-bearing walls of your argument may be unsound. Focus on the claims marked fatal before proceeding.

A cracked verdict is not a failure — it is the system working as intended. Better to discover the cracks now, in private, than in peer review.

CHAPTER IX

Notes, History & Export

A research diary, a complete archive, and everything in portable formats.

Science is not just the final result — it is the messy trail of thoughts that got you there. The Notes tab gives you a place to jot down observations, hypotheses, and methodological notes as you work through your sources. They are timestamped and searchable, and they persist between sessions.

History

Every synthesis, deep dive, citation trace, and storm is automatically saved to your analysis history. The History tab shows them all in reverse chronological order — you can expand any entry to see the full output, copy it, or export it as Markdown. The history holds up to 50 entries per analysis type before older ones are pruned.

Export and Import

The Export button in the top bar saves your entire ResearchForge state — sources, lenses, notes, history, settings, everything — as a single JSON file. You can import it later on a different machine, share it with a collaborator, or keep versioned backups. Your research state is your data; it should always be portable.

Individual analyses can be exported as Markdown (.md) files directly from their output panels. These are formatted for readability and include metadata (model used, lenses active, timestamp).

Keyboard shortcuts. Ctrl/Cmd + S saves state immediately. Ctrl/Cmd + E exports the full state as JSON. These work from any tab.

CHAPTER X

Settings & Backends

Wiring up the intelligence. Local models, cloud APIs, and the fallback chain that keeps the forge running when things go sideways.

The Settings tab is where you configure your AI backends, adjust visual preferences, and monitor connection health. It is intentionally minimal — the forge should not require an engineering degree to configure.

Backend Configuration

Click any backend to select it as primary. Each has its own configuration panel:

Local backends (Ollama, WebAI, StudioLM) need a server URL, a chat model name, and an embedding model name. The default chat model is llama3.2; the default embedding model is nomic-embed-text. WebAI and StudioLM also let you choose between Ollama-compatible and OpenAI-compatible API formats.

Anthropic needs your API key and model name (default: claude-sonnet-4-20250514). Note that Anthropic does not provide embeddings — keep a local backend running for indexing.

The Fallback Chain

When fallback is enabled (the default), ResearchForge tries each configured backend in sequence if the primary fails. The default order is Ollama → WebAI → StudioLM → Anthropic, but you can configure multiple backends so that a network hiccup does not interrupt a long synthesis run.

The "Check All Backends" button probes every configured backend and updates their status indicators. Green means connected. Red means offline. Grey means unconfigured.

Visual Settings

CHAPTER XI

Workflows & Recipes

Three practical workflows, step by step. Not rules — just patterns that have worked for people like you.

Literature Review Workflow

Start by uploading all your papers to the Source Vault. Let the filename heuristics do the sorting, then adjust any misplaced files. Index everything. Load the Epistemic Rigor preset. Run a broad synthesis: *"What is the current consensus on [your topic], and where does it break down?"*

Read the output for contradictions and gaps. Use the Citation Map to trace the most contested claims. Deep Dive into the methodological questions that concern you. Run a Storm on your emerging thesis. Take notes throughout. Export everything.

This workflow typically takes two to four hours for a 20-paper collection. It does not replace reading the papers — it shows you which parts of which papers to read most carefully.

Hypothesis Pressure-Testing

Upload the evidence that supports your hypothesis and the evidence that challenges it. Use the trust weights to mark your confidence level in each source. Run a synthesis asking the forge to pressure-test the hypothesis. Then Storm it at Aggressive intensity.

If it survives, you have grounds for confidence. If it cracks, the Storm output tells you exactly where — and that is more valuable than the confidence would have been.

Grant Preparation

Upload your preliminary data and the key papers you plan to cite. Load the Funding Feasibility preset. Run a synthesis asking: *"Evaluate this research proposal's significance, innovation, and feasibility against the uploaded evidence."*

CHAPTER XII

Philosophy of the Forge

Why we built a tool that argues with you. And why that might be the most useful thing a research tool can do.

There is a temptation, when you build tools with language models, to make them agreeable. To have them summarise papers in encouraging tones, generate text that flatters the hypothesis, and avoid saying anything that might cause the user discomfort.

We went the other direction.

ResearchForge is designed to argue with you — politely, rigorously, and with your own sources as ammunition. The Storm does not exist to validate your conclusions; it exists to attack them. The Contradiction Scanner does not exist to reassure you that your literature is consistent; it exists to find the places where it is not. The Synthesis Engine does not summarise; it interrogates.

This is not contrarianism for its own sake. It is the belief — held with some conviction — that the best research tools are the ones that help you become more wrong more quickly. Because every false certainty you catch in private is a retraction you never have to write.

The Limits of AI-Assisted Research

A language model, however sophisticated, does not understand your field. It does not know which findings will replicate. It cannot tell a rigorous study from a well-written bad one by reading the abstract. It will sometimes flag contradictions that are not contradictions, miss contradictions that are obvious to a domain expert, and occasionally produce analysis that sounds authoritative but is quietly wrong.

This is why every output in ResearchForge is labelled as suggestive, not authoritative. The confidence markers — [STRONG EVIDENCE], [WEAK EVIDENCE], [CONTRADICTION] — are the model's best guesses, not verdicts. Your expertise remains the final arbiter. Always.

Epistemic Humility as a Practice

The most useful thing a research tool can teach you is not what to think, but how to doubt. How to sit with the discomfort of a finding that does not fit. How to resist the urge to smooth over a contradiction in the name of narrative coherence.

ResearchForge is, at its core, a practice in epistemic humility. Use it well — which is to say, use it to find the things you would rather not know.

*The forge does not care about your hypothesis.
It cares about whether your hypothesis deserves to survive.*

APPENDIX A

Keyboard Shortcuts & Quick Reference

Shortcut	Action
Ctrl/Cmd + S	Save state immediately
Ctrl/Cmd + E	Export full state as JSON
Theme toggle (top bar)	Switch between dark and light mode
Sidebar search	Filter research threads by name or content

Supported File Formats

The Source Vault accepts **.txt**, **.md**, and **.pdf** files. PDF text is extracted via pdf.js (if available in the browser) with a raw byte-scan fallback. For best results with scanned PDFs, convert to text before uploading.

Default Models

Purpose	Default Model	Notes
Chat / Synthesis	llama3.2	Configurable per backend
Embeddings	nomic-embed-text	Required for RAG indexing
Anthropic Chat	claude-sonnet-4-20250514	Requires API key

Complete Lens Reference

Epistemic Rigor

Methodology: *Is the methodology appropriate? Where are the design flaws, confounds, and threats to validity?*

Evidence Quality: *How strong is the evidence? What is the effect size, confidence interval, and statistical power?*

Causal Inference: *Is the causal claim justified? What alternative explanations exist?*

Falsifiability: *What would falsify this claim? Has the theory been tested against counter-evidence?*

Internal Consistency: *Do the claims cohere? Where do results contradict the framework?*

Reproducibility

Replication: *Could this study be replicated? What information is missing?*

Data Access: *Is the data available for independent verification?*

Code & Analysis: *Is the analysis pipeline reproducible?*

Materials: *Are the materials, instruments, or reagents fully specified?*

Reporting: *Does the reporting follow field-specific guidelines?*

Interdisciplinary

Bridges: *What concepts from adjacent fields illuminate this work?*

Translation: *How well does this finding translate across domains?*

Methods Import: *What techniques from other fields could strengthen this?*

Theory Integration: *Can competing frameworks be integrated?*

Broader Impact: *What are the societal and ethical implications?*

Funding Feasibility

Significance: *Does this address an important problem?*

Innovation: *What is genuinely novel here?*

Approach: *Is the methodology rigorous and feasible?*

Investigator: *Does the team have the expertise?*

Budget & Timeline: *Is the budget justified and competitive?*

Meta-Analysis

Search Strategy: *Is the literature search comprehensive?*

Heterogeneity: *How much do included studies vary?*

Publication Bias: *Is the evidence base distorted by selective publication?*

Effect Estimation: *Is the pooled effect estimate trustworthy?*

Certainty of Evidence: *What is the overall certainty via GRADE?*

Troubleshooting

Backend shows red dot

Your primary AI backend is not responding. Check that the server is running (for Ollama: run `ollama serve` in your terminal). Verify the URL and port in Settings. Click "Check All Backends" to see which are available.

Indexing fails

Indexing requires a local backend with embedding support. If you are using Anthropic as your primary backend, ensure Ollama or another local server is also running — the system will use it for embeddings via the fallback chain.

PDF text extraction is empty or garbled

Some PDFs contain only images (scanned documents). ResearchForge uses pdf.js for text extraction, with a raw byte-scan fallback. For scanned PDFs, convert to text using OCR software before uploading.

Synthesis output is generic or ignores sources

Make sure your sources are **indexed** — check the Source Stats bar for the indexed percentage. Without indexing, the RAG system cannot retrieve relevant chunks from your materials. Also ensure at least one lens is active in the Lenses tab.

Storage quota exceeded

ResearchForge stores state in browser localStorage. If you have large source collections, this can fill up. Export your state, clear the browser storage, and re-import. Consider splitting very large projects across multiple threads.

Dark mode text is hard to read on some screens

APPENDIX D

Glossary

Chunk — A segment of roughly 800 characters extracted from a source document. Chunks overlap by 200 characters to preserve context across paragraph boundaries.

Deep Dive — A focused analysis using a single lens and focus sub-topic, operating at specialist-reviewer depth.

Embedding — A vector representation of text that enables semantic similarity search. Generated by models like nomic-embed-text.

Fallback Chain — A sequence of configured backends that ResearchForge tries in order if the primary backend fails.

Focus Pill — A predefined sub-topic within a lens that narrows the analysis to a specific aspect (e.g., "power-analysis" within the Evidence Quality lens).

GRADE — Grading of Recommendations, Assessment, Development, and Evaluations — a framework for rating the certainty of evidence in systematic reviews.

Lens — A mode of interrogation — a core question and weight that determines how the synthesis engine examines your research materials.

Preset — A pre-configured set of five lenses designed for a specific type of analysis (e.g., Epistemic Rigor, Reproducibility).

Quick Scan — A client-side contradiction detector that uses TF-IDF cosine similarity and negation analysis. Works without any AI backend.

RAG — Retrieval-Augmented Generation — the technique of searching indexed sources for relevant chunks and including them as context in AI prompts.

Source Vault — The upload and management area for research materials, organised into three trust-tier folders.

Storm — An adversarial pressure-testing tool that runs three sequential passes of increasing intensity against your conclusions.

Synthesis Engine — The core analysis tool that reads sources through active lenses and produces structured output including contradictions, citation chains, and gaps.

*This manual was set in Crimson Text, designed by Sebastian Kosch.
Interface monospace: IBM Plex Mono, designed by Mike Abbink for IBM.*

ResearchForge — Synthesise & Pressure-Test
Version 1.0 · 2026

