

THE GENERAL RESEARCH LAB

Offline Research Platform

User Manual & Onboarding Guide

Field: General Research · \ Offline.Ltd

8 Tabs	30+ Tools	100% Offline	RAG Local AI
------------------	---------------------	------------------------	------------------------

This manual covers every tool, feature, and workflow in The General Research Lab — a complete offline research platform that runs entirely in your browser with local AI via Ollama. From document indexing and RAG-powered research to wiki building, concept mapping, Kanban task tracking, and agentic autopilot research loops, this guide will help you get the most from every feature.

Version 1.4 · For use with The_General_Research_Lab_V_1_4.html

Table of Contents

1 Getting Started

System requirements, opening the file, and first steps

2 Interface Overview

Header, tab bar, content area, and navigation

3 Saving & File Management

Auto-save, export, import, and state management

4 Research Tools

Document Library, RAG Engine, Summarizer, Comparator, Insights, Autopilot

5 Knowledge Tools

Wiki Builder, Quick Notes, Glossary, Journal, Concept Map

6 Analysis & Synthesis

Timeline, Source Matrix, Decision Matrix, Fact-Check, Word Cloud, Bias Analyzer

7 Organization & Planning

Task Tracker, Dashboard, Reading List, Outline Builder, Unified Search

8 Creative & Writing

Draft Builder, Headline Analyzer, Style Checker, Idea Generator, Argument Builder

9 Utilities

Calculator, Date Calculator, File Statistics, Export Center

10 Collaboration

Package Creator, State Merger, Side-by-Side Comparator

11 Knowledge Nexus

Auto-growing knowledge graph with force-directed layout

12 Keyboard Shortcuts

Quick navigation and save shortcuts

13 Troubleshooting

Common issues and solutions

14 Accuracy & Limitations

What to expect and where to be careful

15 Disclaimer

Legal notices, data security, and AI content advisory

16 Glossary

Key terms and definitions

1. Getting Started

The General Research Lab is a single HTML file that runs entirely in your browser. No installation, no server, no internet connection required for core features. AI-powered features require Ollama running locally.

System Requirements

Requirement	Details
Browser	Chrome 90+, Firefox 88+, Safari 15+, Edge 90+
JavaScript	Must be enabled
Ollama	Required for AI features (RAG, summaries, analysis)
Embedding Model	nomic-embed-text (run: ollama pull nomic-embed-text)
Chat Model	Any Ollama model (default: llama3.2)
Storage	localStorage for state, IndexedDB for embeddings

How to Open

For basic features (wiki, notes, tasks, outline), simply double-click the HTML file. For AI-powered features (RAG, summaries, analysis), serve the file via a local HTTP server to avoid CORS issues:

```
python3 -m http.server 8080
```

Then open `http://localhost:8080` in your browser.



PRO TIP

Email the HTML file to yourself so you always have it on every device. It works on phones and tablets too — the interface is fully responsive.

2. Interface Overview

The interface is organized into a header bar, a tab bar with eight main tabs and two system tabs, and a content area. Each tab contains multiple sub-tools accessible via clickable cards.

The Header

The header shows the red backslash (\) brand mark, the product name, the Ollama connection status indicator (click to refresh), the field category, and a link to Offline.Ltd. The Ollama status shows green when connected, red when offline, and amber while checking.

The Tab Bar

Tab	Contents
RESEARCH	Document Library, RAG Engine, Summarizer, Comparator, Insights, Autopilot
KNOWLEDGE	Wiki Builder, Quick Notes, Glossary, Journal, Concept Map
ANALYSIS	Timeline, Source Matrix, Decision Matrix, Fact-Check, Word Cloud, Bias Analyzer
ORGANIZATION	Task Tracker, Dashboard, Reading List, Outline Builder, Unified Search
CREATIVE	Draft Builder, Headline Analyzer, Style Checker, Idea Generator, Argument Builder
UTILITIES	Calculator, Date Calculator, File Statistics, Export Center
COLLABORATION	Package Creator, State Merger, Side-by-Side Comparator
NEXUS	Living Knowledge Nexus — auto-growing concept graph
FILES	Export/Import lab state, per-tool export, recent exports
HELP	Searchable documentation for every tool

3. Saving & File Management

The lab auto-saves to localStorage on every action. Your data persists between sessions in the same browser. Embeddings are stored separately in IndexedDB for performance.

Exporting State

1. Go to the FILES tab
2. Click EXPORT FULL LAB
3. A JSON file downloads containing your complete lab state
4. Keyboard shortcut: Ctrl+E / Cmd+E

Importing State

1. Go to the FILES tab
2. Click IMPORT FULL LAB
3. Select a previously exported JSON file
4. Keyboard shortcut: Ctrl+I / Cmd+I



CAUTION

Importing a state file overwrites all current data. Export your current state first if you want to keep it.

Export Center

The Export Center (Utilities tab) offers four formats: Markdown, CSV, HTML (printable / save as PDF), and JSON. Select which sections to include (Wiki, Notes, Chat, Tasks, Timeline, Outline) and preview before downloading.

Encrypted Export (AES-256-GCM)

The Files tab offers optional password-based encryption for exports. Enter a password and click any lock button to produce an encrypted `.enc` file. The encryption uses AES-256-GCM with PBKDF2 key derivation (310,000 iterations, SHA-256) — the same standard used by modern password managers. Both full lab exports and individual tool exports can be encrypted.

To import an encrypted file, enter the same password and click IMPORT ENCRYPTED. A wrong password produces a clear error — no partial or corrupted data is imported.



CAUTION

Encrypted export protects the exported file only. Your live lab state in localStorage is not encrypted. Anyone with physical access to your browser can read it. For sensitive research, combine encrypted exports with standard device security practices.



PRO TIP

Name your state files descriptively: `climate_research_march2025.json` is much more useful than `export.json`. Use encrypted export when sharing files or storing them on shared drives.

4. Research Tools

The Research tab is the core of the lab. It contains six tools for document management, AI-powered research, summarization, comparison, insight extraction, and autonomous research loops.

Document Library / Indexer

Upload .txt and .md files to build your research corpus. The indexer uses Ollama's nomic-embed-text model to create vector embeddings for each text chunk, enabling semantic search via RAG.

Feature	Description
Upload	Drag-and-drop or click to browse for .txt/.md files
Indexing	Chunks text (800 chars, 200 overlap), embeds via nomic-embed-text
Evidence Tiers	A (3.0x weight), B (1.5x), C (0.8x) — affects RAG scoring
Clear Embeddings	Remove embeddings while keeping documents
Clear All Docs	Remove everything and start fresh

Research / RAG Engine

The RAG engine combines your indexed documents with local AI for contextual, evidence-grounded answers. It supports multi-turn conversation with context pruning, evidence weighting, and optional second opinion from a different model.

Setting	Description
Primary Model	The Ollama model for main responses (default: llama3.2)
Second Opinion	Enable a second model to compare answers
RAG Retrieval	Toggle document context injection (top 6 chunks)
Evidence Weighting	Multiply similarity scores by tier weight
Conversation Pruning	Auto-prune old messages to reduce bias



PRO TIP

Enable Second Opinion when researching controversial topics. Having two models disagree productively surfaces assumptions you might miss with one model.

Smart Summarizer

Paste any text, select length (50/150/300 words) and style (Bullet Points, Paragraph, Technical), and generate an AI summary via Ollama.

Multi-Document Comparator

Paste two documents to receive an AI-powered comparison identifying similarities, differences, contradictions, and unique points.

Key Insight Extractor

Extracts four categories from any text: Claims (main assertions), Facts (verifiable statements), Contradictions (internal tensions), and Questions (follow-up inquiries).

Goals & Autopilot

Set a high-level research goal and let the autopilot loop through your indexed documents using RAG + weighted scoring. The system runs up to 10 steps, choosing between summarizing, retrieving, flagging contradictions, or concluding. Everything is logged and the process can be paused, resumed, or stopped at any time.



CAUTION

Autopilot requires Ollama to be running and documents to be indexed. Without indexed documents, the autopilot has no evidence to work with.

5. Knowledge Tools

Wiki Builder

Build a personal knowledge base with interconnected pages. Use `[[Page Title]]` syntax to create links between pages. Missing links auto-create new pages when clicked. Backlinks show which pages reference the current page.

Quick Notes

Fast capture with comma-separated tags. Notes persist across sessions and are searchable via Unified Search.

Personal Glossary

Build a terminology dictionary for your research domain. Terms and definitions are saved to your lab state.

Daily Notes / Journal

Auto-dated journal entries with mood tracking (Focused, Productive, Neutral, Distracted, Tired). Journal content is extracted to the Nexus graph for concept discovery.

Concept Map Builder

Interactive drag-and-drop mind map. The first node becomes the central concept. Click one node then another to toggle connections. Hover to delete. All positions are saved.

6. Analysis & Synthesis

Timeline Builder

Create investigation timelines with dates, descriptions, and categories (Research, Discovery, Analysis, Other).

Literature / Source Matrix

Organize sources with title, author/publication, and key findings in a structured table format.

Pros & Cons / Decision Matrix

Weigh options systematically by listing advantages and disadvantages for any decision.

Fact-Check & Claim Tracker

Log claims with source and verification status (Unverified, Verified, Disputed, False).

Word Cloud & Text Statistics

Local text analysis: word frequency, lexical density, top-30 word bar chart, sentence count, and average word length. No AI required.

Bias / Framing Analyzer

AI-powered analysis identifying loaded language, framing techniques, omissions, source bias, and overall objectivity rating.

7. Organization & Planning

Task Tracker (Kanban)

Drag-and-drop Kanban board with three columns: To Do, In Progress, Done. Tasks persist across sessions.

Project Dashboard

One-screen overview showing document counts, tier distribution, task status, knowledge stats, recent queries, and active tasks.

Reading List

Track books and papers with title, author, and status (To Read, Reading, Read). Change status with inline dropdowns.

Outline Builder

Hierarchical outline with drag-to-reorder, Tab/Shift+Tab indentation (up to 4 levels), and RAG content pull for any item.

Unified Search

Search across wiki pages, notes, chat history, documents, tasks, and outline items simultaneously.

8. Creative & Writing Support

Draft Builder

Start a new draft with title and content. Drafts auto-save to lab state.

Headline / Title Analyzer

AI-powered headline analysis with strength score (1-10), clarity, engagement, SEO assessment, and 5 alternative suggestions.

Style & Consistency Checker

Scans text for tone consistency, clichés, passive voice, word repetition, sentence variety, and readability level.

Idea & Hypothesis Generator

AI-powered brainstorming: 5 research questions, 3 testable hypotheses, 2 unconventional angles, and 1 interdisciplinary connection. Uses RAG context when available.

Argument Builder

Structures logical arguments with refined thesis, 3-5 supporting arguments, counterarguments with rebuttals, logical structure identification, and weakness analysis. Pulls evidence from indexed documents via RAG.

9. Utilities

Calculator & Unit Converter

Convert between Meters, Feet, Kilometers, and Miles with full precision output.

Date & Timeline Calculator

Calculate duration between two dates in days, weeks, months, years, and hours. Shows midpoint date.

File Statistics & Duplicate Finder

Overview of indexed files: total documents, chunks, characters, tier distribution, per-file details, and Jaccard-similarity duplicate detection.

Export Center

Export selected sections (Wiki, Notes, Chat, Tasks, Timeline, Outline) in Markdown, CSV, HTML, or JSON format. Live preview before download.

10. Collaboration

Collab Package Creator

Export a collaboration package containing wiki, notes, chat history, tasks, outline, document metadata, and nexus graph as a single JSON file.

Multi-State Merger

Intelligently merge multiple JSON states. Deduplicates wiki pages and tasks, combines chat histories, outlines, and nexus graphs.

Side-by-Side Comparator

Guide for running two instances simultaneously using incognito windows, different browsers, or duplicate files for parallel research approaches.

11. Knowledge Nexus

The Living Knowledge Nexus is an auto-growing knowledge graph that extracts concepts and relationships from every action: indexing documents, saving wiki pages, writing journal entries, and asking research queries. Each extraction sends a prompt to Ollama to identify 2-4 key concepts, which become nodes in a force-directed graph.

Feature	Description
Auto-growth	Concepts extracted automatically from documents, wiki, journal, and queries
Evidence Tiers	Nodes colored by tier: red (A), amber (B), steel (C), grey (manual)
Force Layout	30-iteration physics simulation for readable positioning
Drag & Drop	Reposition nodes manually; edges redraw in real time
Click to Query	Click any node to jump to RAG Engine with a pre-filled query
Manual Add	Type concepts directly; auto-connected to nearest existing node

12. Keyboard Shortcuts

Ctrl+S / Cmd+S	Save state to localStorage
Ctrl+E / Cmd+E	Export full lab as JSON
Ctrl+I / Cmd+I	Import lab state from JSON
Enter	Send message in RAG chat, add items in wiki/kanban/outline/concept/nexus
Tab / Shift+Tab	Indent / outdent in Outline Builder
Enter / Space	Activate tool cards (accessibility)

13. Troubleshooting

Ollama shows 'offline'

Ensure Ollama is running (`ollama serve`). If using `file://` protocol, switch to HTTP server to avoid CORS.

RAG returns no context

Documents must be indexed first. Click **INDEX ALL DOCUMENTS** in the Document Library. Ensure `omic-embed-text` is pulled.

localStorage cleared

Browser privacy modes or storage limits can clear data. Export regularly as JSON backup. Embeddings survive in IndexedDB but the state itself needs localStorage.

Import fails

Ensure the JSON file is valid. Open it in a text editor to check for corruption. The file must be a valid lab state export.

Slow indexing

Indexing calls Ollama for every chunk. Yield happens every 3 chunks to keep the UI responsive. Large document sets may take several minutes.

Autopilot stops unexpectedly

Autopilot has a 10-step maximum. It also stops if the model's response suggests the goal has been met. Check the log for details.

14. Accuracy & Limitations

Feature	Notes
RAG Retrieval	Quality depends on embedding model and document chunking. Top-6 chunks may not capture all relevant context.
AI Summaries	Generated by local LLM. May contain inaccuracies or hallucinations. Always verify against source material.
Evidence Weighting	Tier weights (3.0/1.5/0.8) are heuristic. They influence ranking but don't guarantee accuracy.
Bias Analyzer	AI-based framing detection is suggestive, not definitive. Use as a starting point for critical reading.
Autopilot	Autonomous but bounded (10 steps). May pursue tangential paths. Review logs and override as needed.
Word Cloud	Stop word list is English-only. Frequency counts exclude common words but may miss domain-specific stop words.



CAUTION

All AI-generated content should be treated as suggestions, not authoritative answers. The General Research Lab is a research tool, not a source of truth. Always verify important claims against primary sources.

15. Disclaimer

No Warranty

The General Research Lab is provided "as is" without warranty of any kind, express or implied, including but not limited to the warranties of merchantability, fitness for a particular purpose, and non-infringement. In no event shall Offline.Ltd be liable for any damages arising from the use of this product.

AI-Generated Content

All AI-powered features rely on local large language models running via Ollama. LLM output may contain hallucinations (plausible but incorrect statements), omissions, systematic bias from training data, and inconsistency across sessions. AI output is suggestive, not authoritative. Always verify claims against primary sources.

Not Professional Advice

This product is a personal research tool. It does not provide and should not be relied upon as a substitute for legal, financial, medical, academic peer review, engineering, safety, or regulatory compliance advice. Consult qualified professionals for decisions with significant real-world consequences.

Data Storage & Security

All data is stored locally in your browser (localStorage for state, IndexedDB for embeddings). No data is transmitted to any server. However, localStorage is not encrypted — anyone with physical access to your device can read it. The optional encrypted export (AES-256-GCM) protects exported files only. Ollama API calls travel over localhost and never leave your machine.

RAG Retrieval Accuracy

RAG depends on vector similarity search, which is probabilistic. Relevant chunks may be missed and irrelevant chunks may be surfaced. Evidence tier weighting influences ranking but does not guarantee quality or correctness of retrieved content.

Third-Party Dependencies

AI features require Ollama, a third-party local model runner. The availability, accuracy, and safety of Ollama and the models it hosts are outside the control of Offline.Ltd. Model behaviour may change between versions.

Intellectual Property & Governing Law

The General Research Lab and all accompanying materials are the intellectual property of Offline.Ltd. Use is governed by LICENSE.txt. This product and its license are governed by the laws of the Netherlands, with disputes subject to the courts of The Hague. Contact: hello@offline.ltd

16. Glossary

RAG	Retrieval-Augmented Generation. Enhances AI responses by injecting relevant document chunks as context.
Embedding	A numerical vector representation of text, used for semantic similarity search.
nomic-embed-text	The default embedding model used by the lab, run locally via Ollama.
Ollama	A local AI model runner. Hosts both embedding and chat models on your machine.
Chunk	A segment of a document (default: 800 characters with 200-character overlap) used for indexing.
Evidence Tier	A quality rating (A/B/C) assigned to documents that affects their weight in RAG retrieval.
Cosine Similarity	A mathematical measure of how similar two embedding vectors are, ranging from -1 to 1.
Pruning	Removing older conversation messages to reduce context bias and save token budget.
Second Opinion	A feature that sends queries to two different models for comparison.
Nexus	The auto-growing knowledge graph that connects concepts extracted from your research.
Kanban	A task management method using columns (To Do, In Progress, Done) for visual workflow tracking.
Backlinks	Wiki pages that reference the current page, enabling bidirectional navigation.
Autopilot	An agentic research loop that autonomously queries your documents toward a stated goal.
IndexedDB	A browser database used to store large embedding vectors separately from localStorage.
CORS	Cross-Origin Resource Security. Browser policy that blocks file:// pages from calling localhost APIs.
Force-Directed Layout	A physics-based graph layout algorithm that uses repulsion and attraction forces.
Flowable	A ReportLab concept: any element (paragraph, table, image) that flows into a document layout.
State File	A JSON export of the entire lab state, importable to restore all data.
Lexical Density	The ratio of unique content words to total words, indicating text complexity.

THE GENERAL RESEARCH LAB

Offline Research Platform

Stay curious.

Version 1.4 · \ Offline.Ltd · offline.ltd