

ParancU

Local-first evidence retrieval for private documents

The problem

Document question answering often fails before generation begins: the system retrieves text that sounds related but does not actually contain the evidence needed to answer.

ParancU focuses on that earlier layer. It retrieves, ranks, and exposes supporting passages so users can inspect the evidence directly.

The ParancU approach

- 1 Load, paste, or scan document text.
- 2 Prepare a searchable corpus locally on the device.
- 3 Embed queries and guiding questions with a multilingual E5 model.
- 4 Combine semantic similarity with lexical and anchor evidence.
- 5 Return the strongest passages with visible evidence and retrieval metadata.

Why local-first matters

- Query embeddings, scoring, ranking, and evidence inspection run on-device.
- The core retrieval workflow does not require an OpenAI key.
- Private source text can remain within the mobile workflow.
- Results remain inspectable instead of being hidden behind a fluent answer.

What the app includes

- Ask: query the currently prepared corpus.
- Upload: import text or PDF content and prepare it locally.
- Scan: capture or select images, run OCR, inspect extracted text, and prepare a corpus.
- Settings: select language and adjust the semantic-lexical retrieval balance.

Research and evaluation

ParancU v1.1 develops the evidence-first retrieval work into a local-first mobile prototype. The current white paper documents the architecture, multilingual E5-small runtime, semantic-lexical score fusion, and retrieval evaluation across three example corpora.

The reported local Top-5 evidence-containment results are 85% for Question Answering, 82% for Sicilia Normanna, and 78% for Multiple Sclerosis.

Current status

ParancU 0.1.0 is available on the Apple App Store for iOS 17 or later.

[View ParancU on the Apple App Store](#)