MiniMind Mermaid Run 2: epoch checkpoint archive

Lucas van de weerd | Lucasgitguy-co | 7 November 2026 | Version 1.0

Repository: Lucasgitguy-co/minimind-mermaid-run2

Reviewed revision: 1d1ff9a5bef6f00d4b19a8e6b32da50e5490f973

Abstract

This report describes three published epoch checkpoints for text-to-Mermaid generation. It documents the export architecture and selection requirements without inferring a quality ordering from the epoch numbers.

1. Scope and related work

This technical report reconstructs the release from pinned repository artifacts and the publisher's earlier portfolio report. It is an AI-assisted documentation synthesis, not an independent rerun or a peer-reviewed publication. Reported measurements belong to the original release records. [1, 2, 3]

The earlier portfolio report provides the organizing distinction between training, model modification, conversion and packaging. This report applies that distinction to one release and uses current records where later experiments supersede the historical overview.

2. Model and method

The ep1, ep2 and ep3 folders each contain model weights, tokenizer files and configuration. The ep3 export declares Qwen3ForCausalLM, eight hidden layers, hidden size 768, vocabulary size 6,400 and float16 weights. The card records Transformers 4.57.6 and the MiniMind-3 upstream family. There is no included vision projector; the export is a text-generation model. [1, 2, 3]

3. Evidence and results

The three saved states establish the existence of an epoch archive. No recorded validation table or held-out diagram evaluation establishes which checkpoint should be preferred. The ep3 loading example is an explicit subfolder choice, not evidence that ep3 is best. [1, 2, 3]

4. Limitations and intended use

Historical training data, exact split and optimization settings remain incomplete. Project dataset links cannot establish those facts. A useful comparison would evaluate all three checkpoints with identical decoding, diagram parsing and graph-semantic metrics on content excluded from training. Apache-2.0 notices are retained.

5. Reproducibility and conclusion

The source snapshot above fixes what was reviewed. The linked source files and source-manifest.json record the evidence used, including file hashes. Consult the repository's usage instructions and recorded dependencies for execution. No new model training, inference benchmark, or claimed state-of-the-art result was introduced in preparing this report. Evidence status: artifact documentation.

References

[1] Lucasgitguy-co. README.md at reviewed revision.

[2] Lucasgitguy-co. ep3/config.json at reviewed revision.

[3] Lucasgitguy-co. docs/PROVENANCE.md at reviewed revision.

[4] Lucasgitguy-co. Local Model Adaptation Across Vision Diagrams Geometry and Game Decisions. Historical portfolio report, 5 November 2026; its proposed v2 experiment is superseded by later release records.