Clinical Graph-JEPA: Predictive Patient-State Knowledge Graphs for Cognitive Decision Support

Kushagra Yadav[1], Nalin Prabhath[2], Amit Lamba[3], James E. Schrager[2], Goeun Han[2], Yining Mao[2]

[1] New York University, [2] The University of Chicago [3] Mangoes.ai

CONTACT ME

READ THE PAPER

Summary

Graph–JEPA uses localized discharge–note context to recover clinically meaningful relationships missing from the graph–only representation, producing a more coherent patient–state knowledge graph.

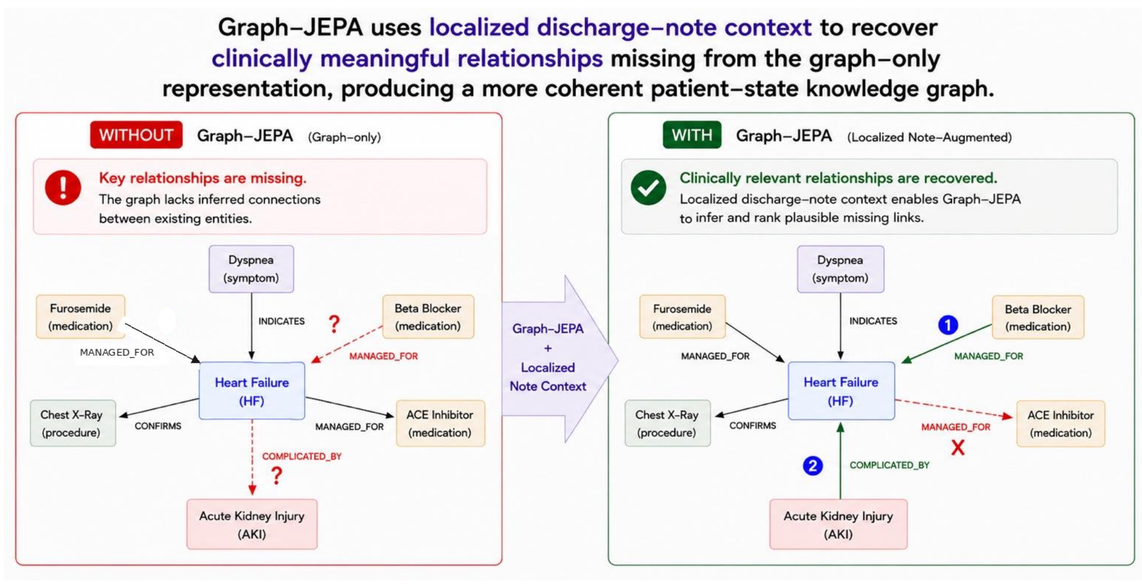

Introduction

Clinical records tell us what happened to a patient, but understanding a patient requires knowing how those facts fit together. We therefore frame the clinical knowledge graph as an **evolving patient-state representation**, where Graph-JEPA learns to predict missing graph regions from the patient context already observed.

Importantly, **more context is not necessarily better context**. We investigate where unstructured clinical information should enter the graph, comparing graph-only, globally injected, and entity-grounded note context. The model acts as a **lightweight refiner** rather than a cold-start predictor, recovering plausible missing relationships from an existing evidence-grounded graph under schema-constrained, leakage-aware evaluation.

Dataset & Cohort

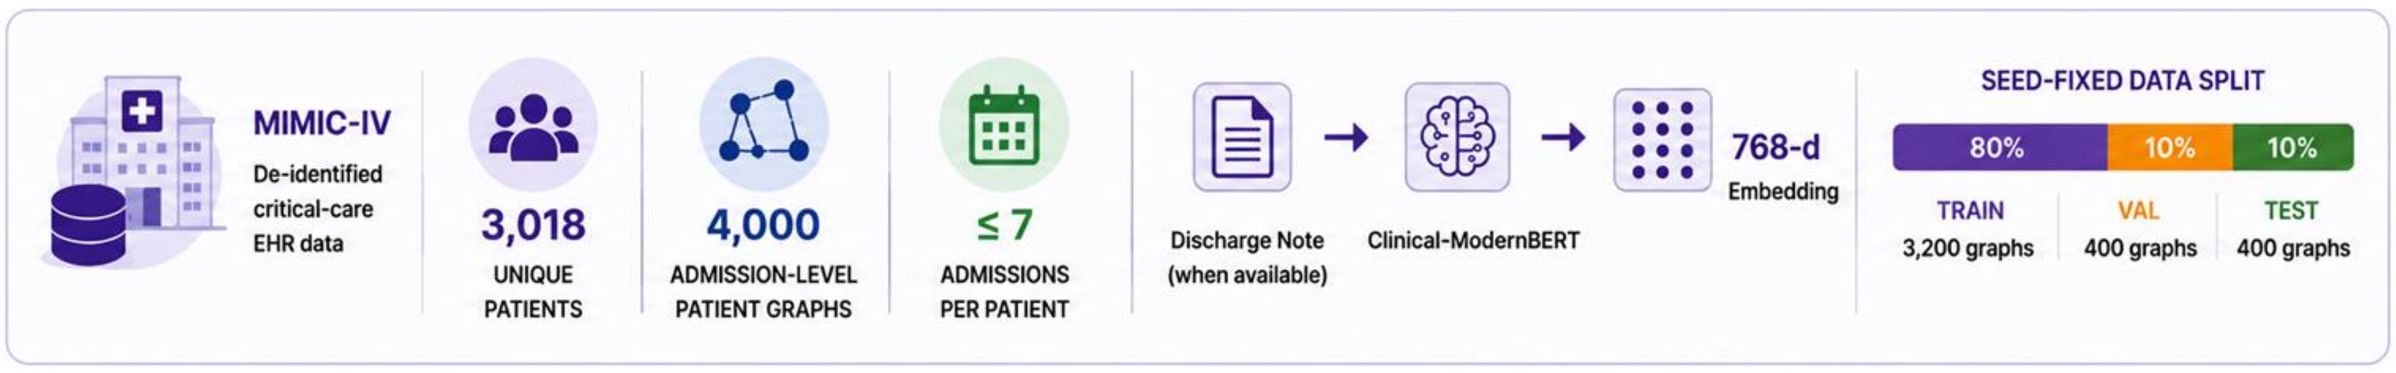

Methodology

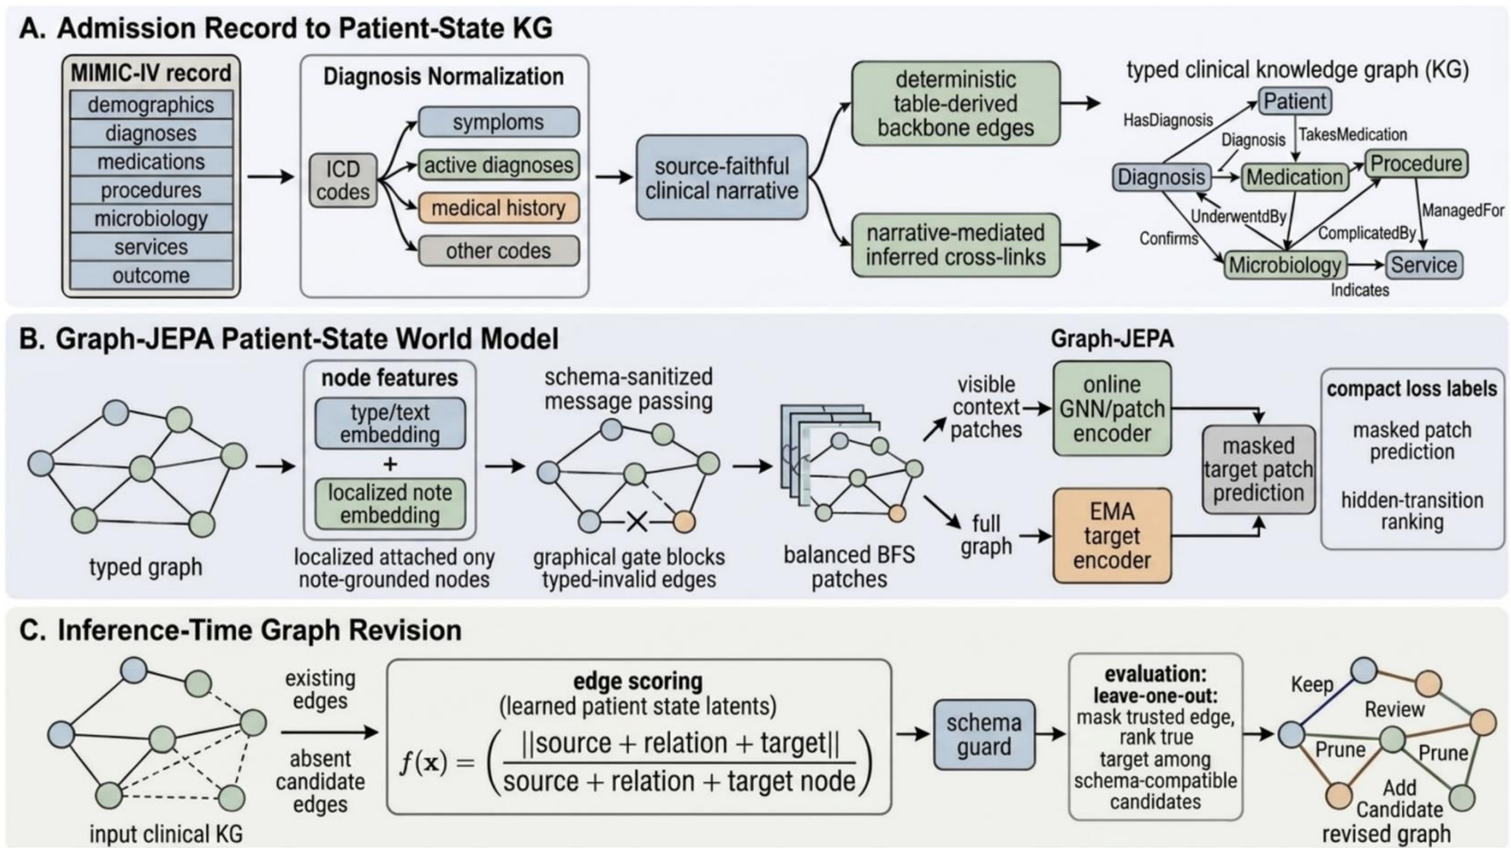

- Patient records become a structured clinical graph.
- Specialist agents infer missing clinical relationships.
- Evidence scoring filters links before integration.

- Parts of the patient graph are masked.
- Visible context predicts the missing patient state.
- The model learns how clinical information fits together.

- Learned representations score missing relationships.
- Clinical schema restricts valid candidates.
- Plausible links are surfaced for graph refinement.

Experimental Setup

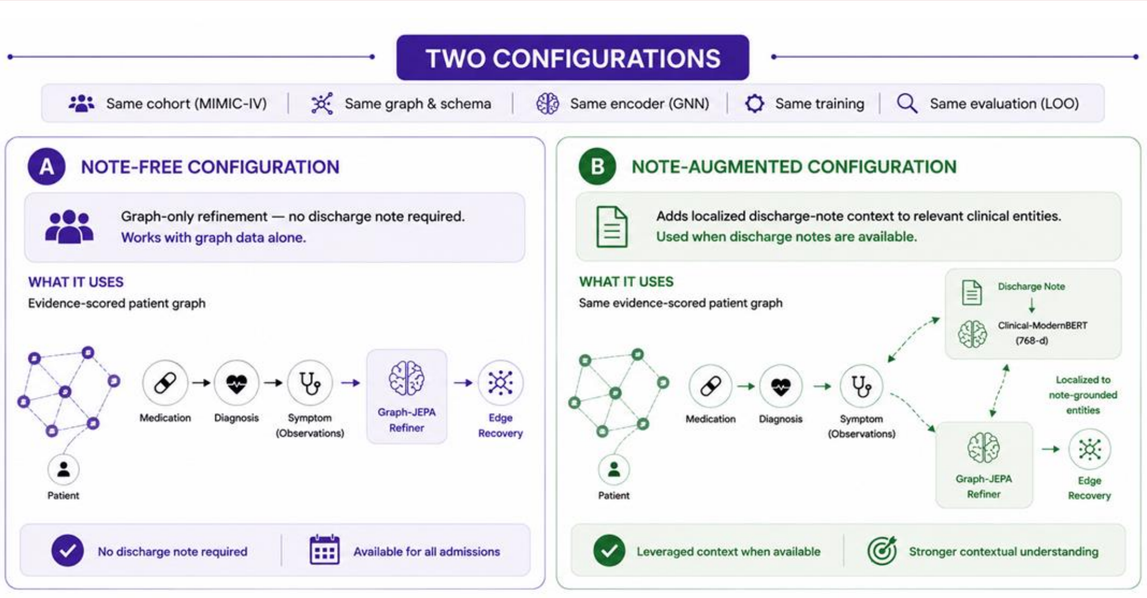

Results

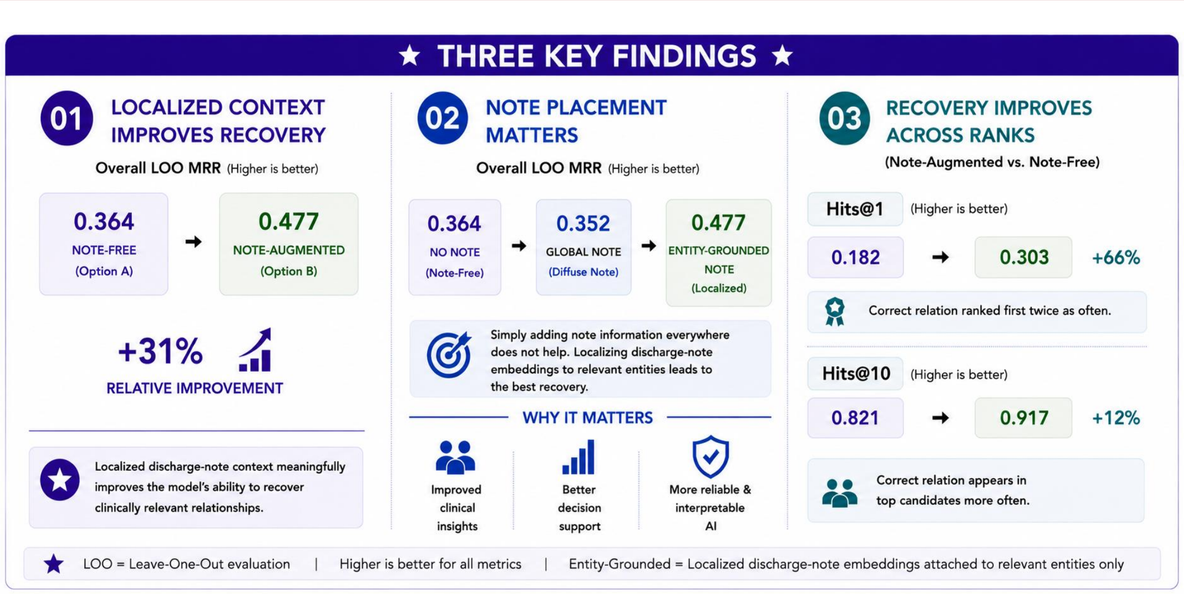

Limitations

- **Single-site cohort:** Evaluation is primarily based on MIMIC-IV
- **Recovery ≠ clinical truth:** Relations are not independently clinician-adjudicated
- **Clinical impact not yet tested:** No prospective evaluation of patient outcomes or decision-support utility.

Future Work

- **External validation:** Clinician-annotated relations across multiple sites.
- **Stronger evaluation:** Patient-disjoint and temporally separated splits.
- **Clinical utility:** Evaluate graph refinement for retrieval and decision-support tasks.

References

[1] Assran et al. (2023). *Self-Supervised Learning from Images with a Joint-Embedding Predictive Architecture.*

[2] Skenderi et al. (2025). Graph-Level Representation Learning with Joint-Embedding Predictive Architectures.

[3] Johnson et al. (2023). *MIMIC-IV, a freely accessible electronic health record dataset.* Scientific Data