

Generative AI and Knowledge Representations

Natural Language Processing-Lab

Indiana University at Bloomington



Large Language Models

- Surprising performance
- Minimize the cost of producing text, images, code, video, and audio

Problems:

- Hallucinations and Trustworthiness
- Bias and Toxicity
- Complete lack of semantic processing or understanding

Solutions:

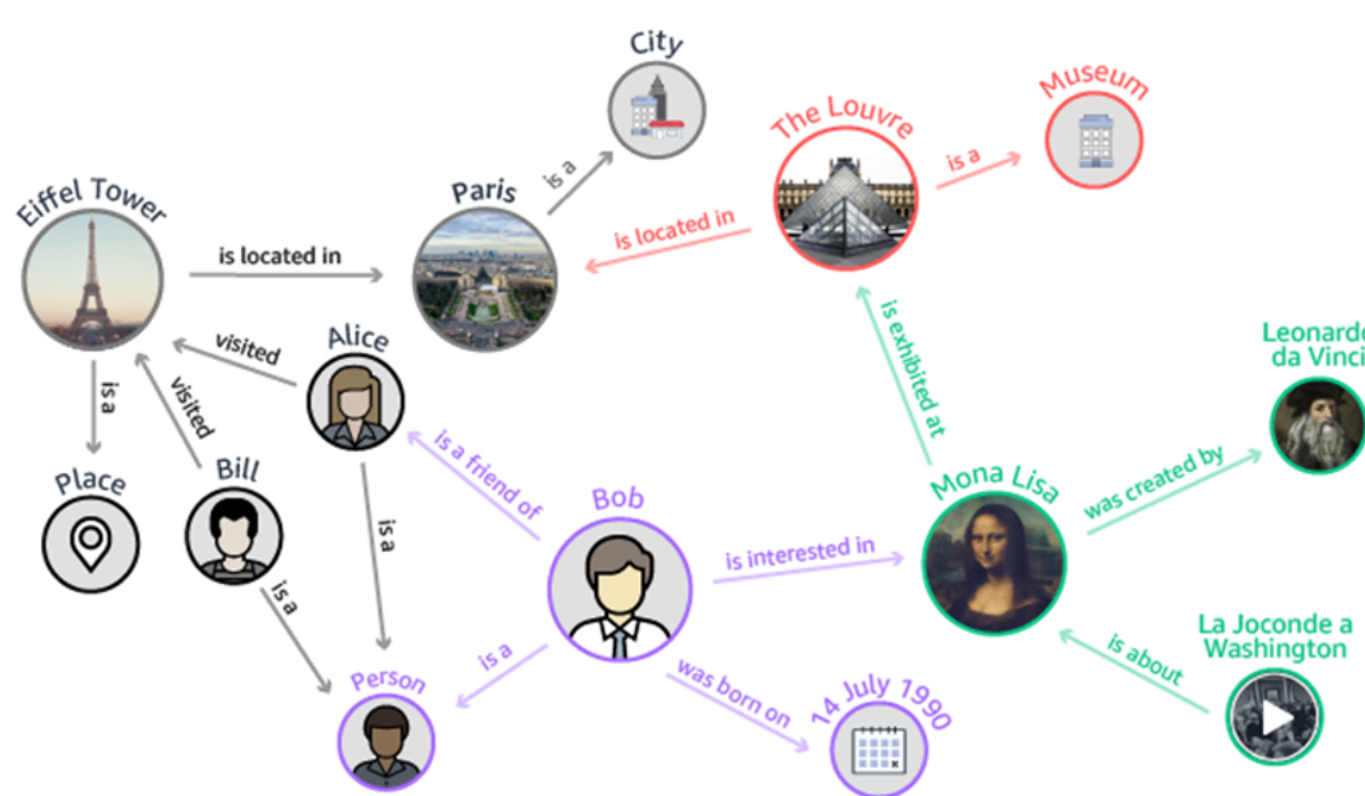
- Training and Tuning using filters for bias, toxicity
- Retrieval Augmentation Generation (RAG) to minimize hallucinations and maximize trustworthiness

Our Solutions:

- Different architecture and data
- Knowledge and Semantics-based

Knowledge Graph Models

Description Logic representation of static knowledge and facts:



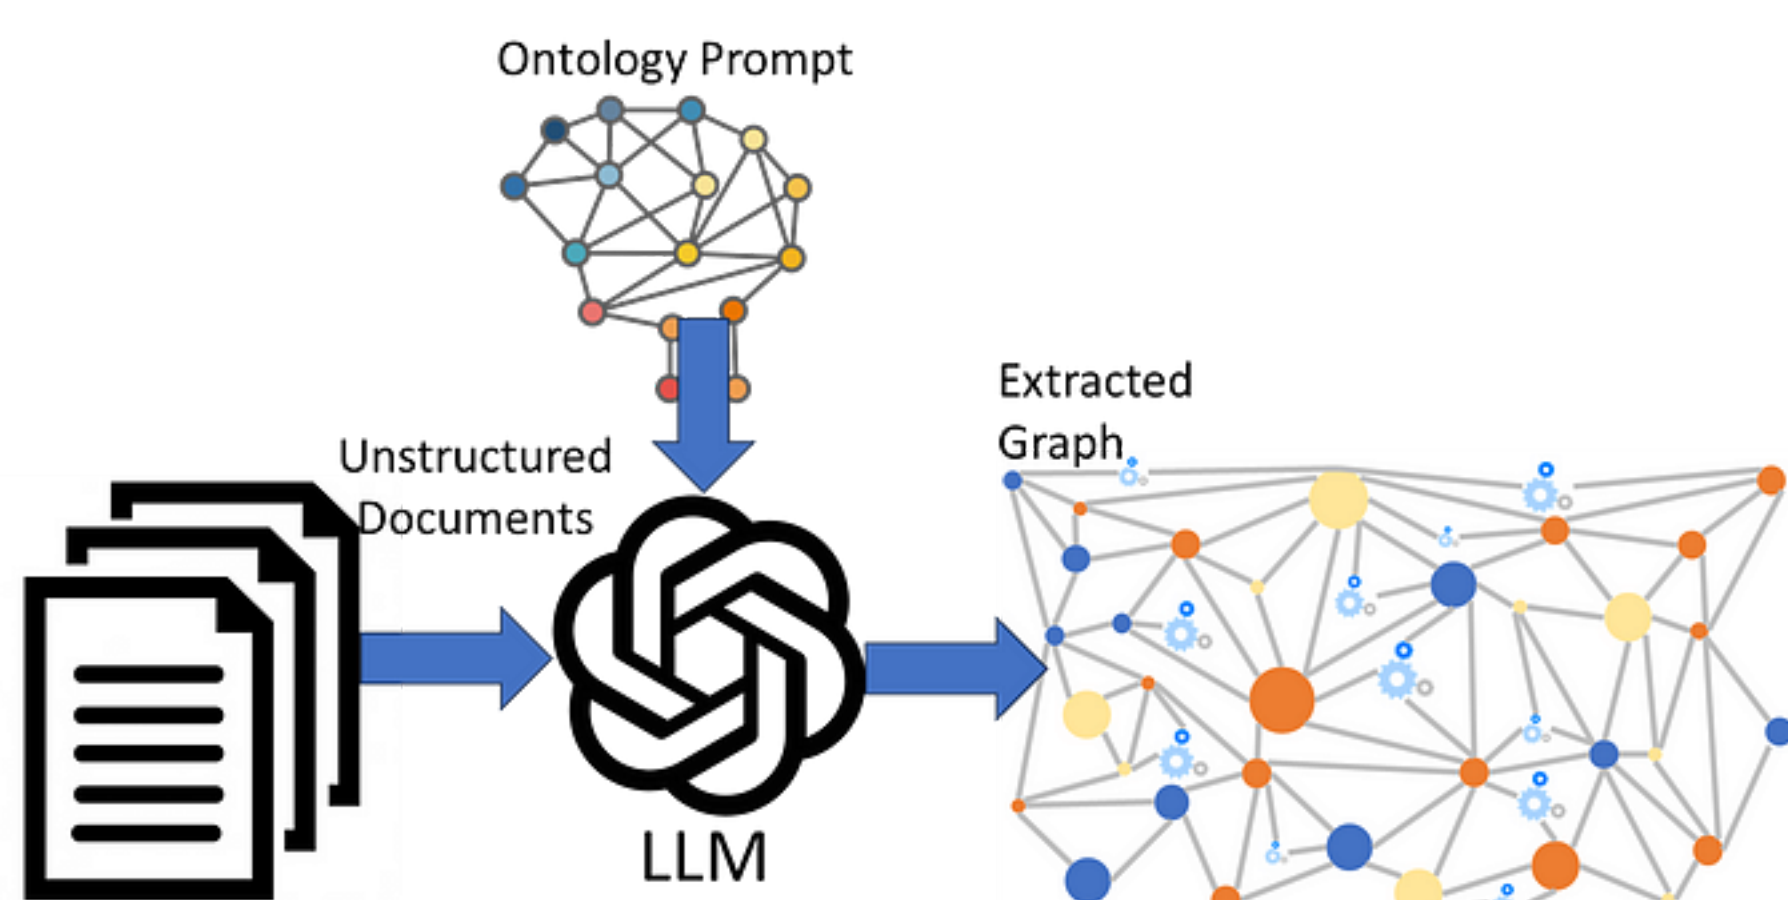
Problems:

- Costly aggregation and maintenance, continuous tracking for updates
- Static knowledge representation
- No encoding of procedural knowledge or event unfolding models

LLMs and Knowledge Representations

Retrieval Augmentation Generation (RAG): not sufficient to eliminate LLM/AGI problems. (See Marcus, 2024)

Automatic unstructured data (text) to Knowledge Graph / Ontology mapping, linking, enabling semantic search and reasoning (see Lawrence, 2024)




Dynamic graphs encoding events and procedural knowledge:


- Graphlets encoding states in events and procedures
- State transition with temporal information (temporal reference, duration)

Event Graphs and Graph Embeddings

Events as graph transformations:



Computing Graph Embeddings over concepts in Knowledge Graphs for semantic AI models:



$$h_v^{(t+1)} = \sigma(W_l \sum_{u \in N(v)} \frac{h_u^{(t)}}{|N(v)|} + B_l h_v^{(t)})$$

Semantic embeddings from KG for Generative AI and NLP applications.

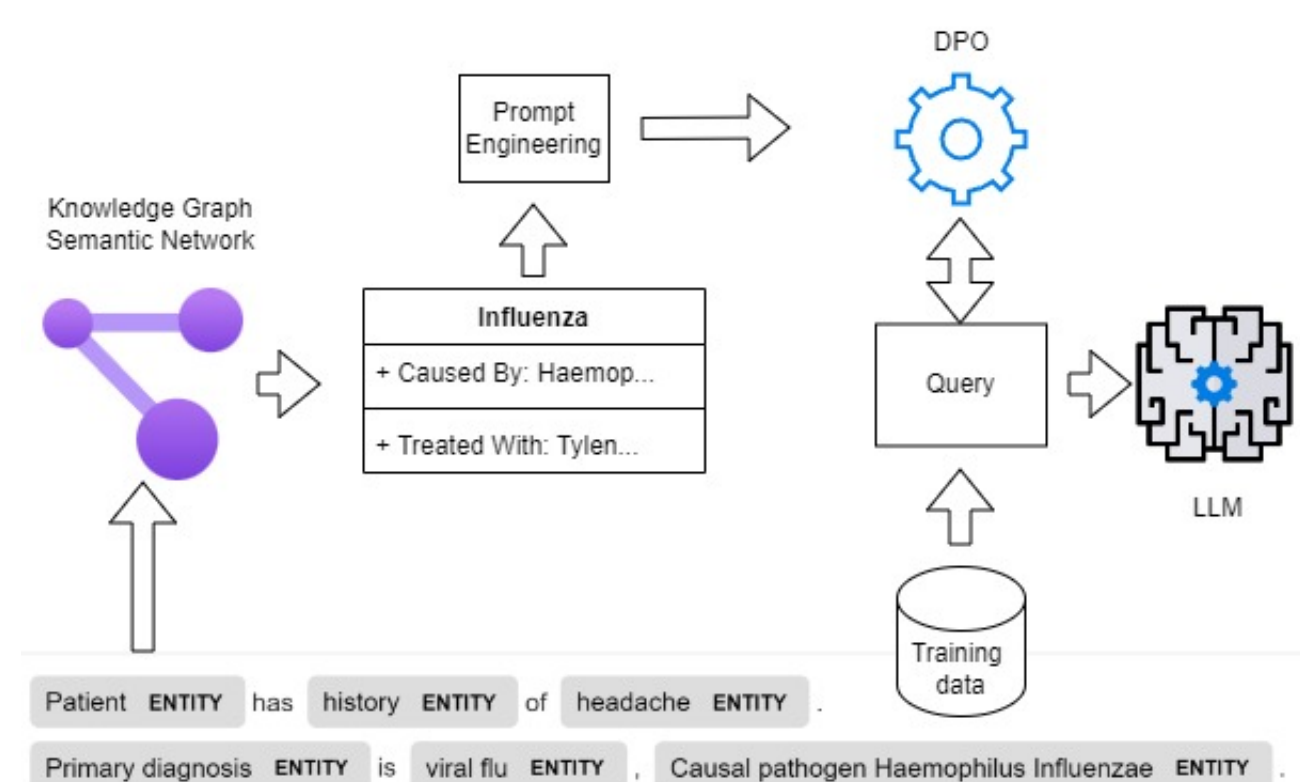
Legal Applications

Compliance Management:

- Tracking changes in regulatory requirements and laws
- Identifying compliance violations or conflicts in products, procedures, services

Medical Applications

- ACL Clinical NLP Subtask on Medical Error Detection and Correction.
- Semantic networks in UMLS ontologies to ground LLMs with structured context and train them with direct preference optimization.



FinTech Applications

- NLP components for financial report processing in Arabic.
- Mapping the trajectory of R&D spending against earnings within SEC reports to assess the direct impact of technology on fiscal performance.
- Evaluating stock response volatility to tech announcements in SEC filings, assessing the market's valuation of technological progress.

Natural Language Processing Lab & Team

The NLP-Lab (<https://nlp-lab.org/>) also working on:

- Natural Language Processing and Large Language Models
- Quantum NLP

