



ONTOLOGIES AND LARGE LANGUAGE MODELS RELATED BUT DIFFERENT

ONTOLOG SUMMIT 2024 FALL SERIES; OCTOBER-NOVEMBER 2023

CHAIRS: MIKE BENNETT AND ANDREA WESTERINEN

Synthesis
15 November 2023

SUMMARY, 1

- LLM use cases/strengths:
 - Q&A and chatbots
 - Information extraction and summarization
 - Translation
 - Classification and clustering
 - Content (visual, audible, text) generation and editing
 - “Copilots” assisting users (initially programmers; being expanded for more general use cases)

SUMMARY, 2

- LLM challenges:
 - Coherence
 - Provenance
 - Semantics
 - Explainability
 - Reasoning
 - ...
- LLM challenges align well with Ontology capabilities => “Hybrid systems”
 - Combining the strengths of LLMs and ontologies/knowledge graphs to overcome weaknesses of each

USAGE OF BOTH LLMS AND ONTOLOGIES, “HYBRID SYSTEMS”

- LLMs ...
 - Extract information from text and aid in its mapping to an ontology
 - Create lists of initial concepts to use when starting or extending an ontology or exemplary KG instances
 - Generate SPARQL queries from natural language
 - Summarize an ontology or KG
 - Aid in alignment of ontologies and KGs
 - Generate competency questions and use cases for an ontology
- Ontologies ...
 - Validate the responses of an LLM
 - Provide training data or prompt inputs to an LLM
 - Support logical reasoning and inference

FALL SERIES

- Discussed “hybrid systems”, provided motivation for developing them, and demonstrated applications/sandboxes based on them
- Highlighted need to keep exploring areas of collaboration, and improving both ontology and LLM development and use
- Various architectures/frameworks show different interactions between ontologies and LLMs
 - Several with explicit feedback loops

SESSIONS AND SPEAKERS

- 4 October 2023, *Kickoff*
- 11 October 2023, *Setting the stage*
 - **Deborah McGuinness**, The Evolving Landscape: Generative AI, Ontologies, and Knowledge Graphs
- 18 and 25 October 2023, *A look across the industry*
 - **Kurt Cagle**, Complementary Thinking: Language Models, Ontologies and Knowledge Graphs
 - **Tony Seale**, How Ontologies Can Unlock the Potential of Large Language Models for Business
 - **Evren Sirin**, Stardog Voicebox: LLM-Powered Question Answering with Knowledge Graphs
 - **Yuan He**, DeepOnto: A Python Package for Ontology Engineering with Deep Learning and Language Models


SESSIONS AND SPEAKERS

- 1 November 2023, *Demos of hybrid systems*
 - **Prasad Yalamanchi**, Harvest Knowledge From Language - Harness the power of Large Language Models and Semantic Technology
 - **Andrea Westerinen**, Populating Knowledge Graphs: The Confluence of Ontology and Large Language Models
- 8 November 2023, *Broader thoughts*
 - **Anatoly Levenchuk**, Hybrid Reasoning, the Scope of Knowledge, and What Is Beyond Ontologies?
 - **John Sowa** and **Arun Majumdar**, Trustworthy Computation: Diagrammatic Reasoning With and About LLMs

DEBORAH MCGUINNESS' EVOLVING LANDSCAPE

- AI is changing business landscape
 - Necessary to understand strengths and weaknesses of LLMs
 - “AI will not replace most knowledge professionals but many knowledge professionals who do not collaborate with generative AI will be replaced”
 - “Generative AI explosion provides ... a unique opportunity to shine and a time to rethink our methods”
- LLMs are “usefully wrong” – providing information to help you think
- Applications of LLMs at RPI
 - Collaborative KG generation by leveraging LLMs for refinement and population (value restrictions and instances) of an existing ontology, in partnership with human
 - Enhancing wine and cheese ontology
 - But could also provide concepts that are a starting point for a new ontology, for human consideration
 - LLM/KG Fact Checker (ChatBS) “sandbox” with questions submitted (multiple times) to OpenAI completion API and entity linking to Wikidata for validation

GARTNER HYPE CYCLE FOR AI, 2023 (FROM D. MCGUINNESS)




KURT CAGLE'S COMPLEMENTARY THINKING

- Mapping LLMs to ontologies/KGs
 - Matching LLM concepts to KG instances over specific classes such as schema.org or NIEM
 - Using a RAG (Retrieval Augmented Generator) plug-in to communicate with an ontology/KG and add to the node-sets or control output transformation
 - Reading Turtle, RDF/XML and JSON-LD
- Mapping ontologies/KGs to LLMs
 - Using URI/IRI references in data and obtaining results with those references
 - Adding KG embeddings (vector space representations) to LLM training




TONY SEALE'S UNLOCKING POTENTIAL

- LLM and ontology “reinforcing feedback loop of continuous improvement”
 - Using ontology/KG to place “guardrails” on LLM outputs
 - Using LLMs to aid in maintenance and extension of ontology
- Information as a continuous stream (~LLMs) or discrete chunks (~KGs)
 - Analogy to System 1 (intuitive/instinctual) and System 2 (reasoning based) thinking



WORKING MEMORY GRAPH (FROM T. SEALE)




This WMG combines internal and external data and concepts, serving as a context for the task at hand. It can be used:

- as context injected into prompts
- to construct queries to the KG (GRAG)
- to convert LLM responses into a graph
- to reason and validate LLM responses
- to generate entirely new prompts
- to assist in orchestration (Graph of Thought)
- to suggest exertions to the ontology


EVREN SIRIN'S LLM-POWERED QUESTION ANSWERING

- Stardog Voicebox combines LLM and graph database technology to:
 - Take a description of ontology and create it
 - Turn natural language query into SPARQL
 - Provide context for decisions and debug/repair queries
- Built on:
 - Open-source foundational model, MPT-30B
 - Fine-tuned with ~20K SPARQL queries
 - Vector embedding and search via MiniLM-L6-v2 language model



YUAN HE'S DEEP ONTOLOGY ENGINEERING

Overview




- DeepOnto
- Python package for ontology engineering with deep learning and LMs

PRASAD YALAMANCHI'S HARVESTING KNOWLEDGE

■ TextDistil


- Inputs – text documents;
Outputs – NQuad files and JSON
- Models trained on domain-specific variables, and training data labeled using taxonomy
- Ontology for organization/semantics (human defined)



TEXTDISTIL (FROM P. YALAMANCHI)

■ TextDistil


- Query in NL parsed to ontology concepts and used to generate query to KG
- Triples returned with provenance from ingested documents
- LLM used to summarize response



ANDREA WESTERINEN'S POPULATING KNOWLEDGE GRAPHS

- Overview of open-source tooling to parse news articles (Deep Narrative Analysis, DNA)
 - Create knowledge stores with data from text stored in RDF graphs
 - Enabling aggregation of textual information within and across documents
 - To efficiently compare and analyze collections of text to understand patterns, trends, ...
- Prompts sent to OpenAI chat completion API for:
 - Narrative analysis
 - Rhetorical devices and viewpoint interpretations
 - Sentence analysis
 - Linguistics (tense, voice, errors, ...), rhetorical devices and mapping to ontology
- LLM JSON responses (already mapped to the ontology) used to generate RDF
 - Which is stored in graph database

DNA ARCHITECTURE (FROM A. WESTERINEN)



ANATOLY LEVENCHUK'S COGNITIVE ARCHITECTURES

- A cognitive system/agent is a *cognitive architecture* with a collection of KGs, LLMs and other knowledge representations
 - Cognitive architecture refers to both a theory about the structure of the human mind and to a computational instantiation of such a theory used in the fields of artificial intelligence (AI) and computational cognitive science (https://en.wikipedia.org/wiki/Cognitive_architecture)
- Where KGs are discriminative declarations of “what is in the world” and LLMs are generative
- Both have roles in knowledge evolution
- “Looking at LLMs as chatbots is the same as looking at early computers as calculators. We're seeing an emergence of a whole new computing paradigm, and it is very early.”

KNOWLEDGE EVOLUTION (FROM A. LEVENCHUK)

- LLM usage with ontologies for knowledge evolution



JOHN SOWA'S/ARUN MAJUMDAR'S TRUSTWORTHY COMPUTING

- Large language models cannot do reasoning, but find and apply reasoning patterns from training data
- Important to note that “thinking in language” is only one form of reasoning
- Systems developed by Permion use LLMs for summarization/synthesis
 - But restrict responses based on the ontology
- Combine LLMs with a “scaffolding model” (vector, matrix and tensor-based) => ontology and methods of diagrammatic reasoning based on conceptual graphs (CGs)
 - Where ontology is derived/tailored to policies, rules, and specifications of the project or business

what is the role of ACE2 and what is the capital of Alaska?

Source relevance score: 0.09

No strict answer found. Try "Broad".
Missing:
• what is the capital of Alaska

Sentence Selection
☒ Strict
☐ Broad

LLM Synthesis
☐ Off
☒ Strict
☐ Broad

Translate

Score Answer Sentences

The role of ACE2 is to serve as a receptor for both SARS-CoV-2 (covid-19) and SARS-CoV-1 viruses. It is found on the viral envelope membrane and the host receptor, specifically in the lung alveolar epithelium and intestinal enterocytes. The binding of the viruses to ACE2 is facilitated by the C-terminal domain. The higher binding affinity of SARS-CoV-2 to ACE2 compared to SARS-CoV-1 is believed to contribute to its greater infectivity. Dimerization of the peptide domain of ACE2 allows for the simultaneous binding of two SARS-CoV-2 viruses, enhancing viral load.



BACKUP SLIDES

WHAT ARE LLMS GOOD AT?

- Excerpted from [blog post by Cobus Greyling](#)
- 1. Proficient in generating text; producing fluent, succinct and precise linguistic expressions
- 2. Demonstrate remarkable proficiency in language comprehension tasks, such as sentiment analysis, text categorization, and processing factual input
- 3. Aptitude for logical and temporal reasoning
- 4. Demonstrate contextual understanding, allowing them to produce coherent responses in line with provided input
 - Challenge is to deliver accurate and highly contextual information at inference time
- 5. Attain impressive results across various natural language processing tasks, encompassing machine translation, text generation, and question answering

WHEN CAN LLMS FAIL?

■ Excerpted from [blog post by Cobus Greyling](#)

1. Display below-average performance and face difficulties in accurately representing human disagreements (NLI - the task of determining whether a given “hypothesis” logically derives from a “premise”)
2. Manifest credibility deficits, giving rise to fabricated information or erroneous facts within dialogues
3. Restricted proficiency in discerning similarity between events
4. Limited ability to perform [abstract reasoning](#), and prone to confusion or errors in complex contexts
 - Mitigated to some degree with innovative prompting techniques
5. Capacity to assimilate, disseminate, and potentially magnify detrimental content found within the acquired training datasets
6. Exhibit social biases and toxicity during the generation process, resulting in the production of biased outputs
7. Limitations in incorporating real-time or dynamic information
8. Sensitive to prompts