ONTOLOGIES AND LARGE LANGUAGE MODELS

RELATED BUT DIFFERENT

ONTOLOG SUMMIT 2024 FALL SERIES; OCTOBER-NOVEMBER 2023

CHAIRS: MIKE BENNETT AND ANDREA WESTERINEN

ABSTRACT, 4 OCTOBER 2023

The opening session of the Ontology Summit 2024 Fall Series overviews the LLM, ontology and knowledge graph landscapes, as well as introducing the participating speakers. The goal of the Series is to understand, discuss and debate the similarities, differences and overlaps across these landscapes.

OVERVIEW

- Ontologies and Large Language Models (LLMs) represent different but related concepts within the fields of artificial intelligence and knowledge representation
  - Ontologies are representations of a knowledge domain
    - Defining the concepts, relationships, properties, axioms and rules within that domain, providing a framework that enables a deep understanding of that subject area
    - Enabling machine reasoning, allowing a system to draw inferences and to derive new information and relationships between entities
    - Knowledge graphs organize and link data as instances of the concepts in the ontologies
  - LLMs are machine learning models
    - Generating human-like responses (including text, images, video and audio) based on inputs ("prompts")
    - Trained on a large corpora of (mostly online) text, learning the patterns and connections between words and images
- The success (and failure) of LLMs has generated much interest in AI and machine learning
  - Leveraged here to discuss the benefits of, and increase awareness of, the value of ontologies and knowledge graphs

GENERATIVE AI LANDSCAPE

- Use cases:
  - 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)
- Generative AI as a New Innovation Platform (Communications of the ACM, Oct 2023)
  - “Bloomberg estimated the market for generative AI hardware and software was … worth $40 billion in 2022 and likely to grow to $1.3 trillion over the next 10 years”
  - “ChatGPT … attracted as many as one billion users as of July 2023 … usage levels seem to be slowing”
  - “At least 335 startups now target generative AI"
  - “Established companies are exploring ways to incorporate generative AI into existing products and services”

GENERATIVE AI CHALLENGES

- Generative AI as a New Innovation Platform (Communications of the ACM, Oct 2023)
  - "Geoffrey Hinton … left his position at Google in May 2023 after warning generative AI would diffuse too much misinformation and become detrimental to society ... especially worried that the current systems had no guardrails limiting bad behavior or social damage"
  - "What to worry about?"
    - "Data privacy, bias, and content ownership issues … the new trillion-dollar question: What is "fair use" of training data for generative AI systems?"
    - "Information accuracy and authenticity" in question ("hallucinations")
    - Need for regulation/self-regulation
    - "Environmental impact ... generative AI's use of computing resources has been increasing exponentially for years, doubling every 6 to 10 months"
    - "Concentration of market power **…** unlikely an open source platform or small players will be able to compete long-term"
    - "Unintended consequences … many occupations … may find their jobs replaced, enhanced, or greatly altered"

DISCUSSION FOCUS

- **Semantic Understanding and Knowledge Representation**: Both ontologies and LLMs represent "knowledge"
  - Ontologies explicitly capture data's semantics regarding entities and their relationships, and are based on a formal, structured, axiom-based/logical representation of a domain of knowledge
  - LLMs, in contrast, have implicit, probabilistic representations of "knowledge", based on the patterns in their training data
- **Technical Assistance and Hybrid Systems**: Both ontologies and LLMs can be used to build technical assistants such as question-answering systems, information extraction systems, chatbots, etc.
  - Used together, they can improve functionality and correctness

SEMANTIC UNDERSTANDING

- Semantics is about “meaning”
  - Agreeing on meaning/interpretation => ”Understanding”
  - Having no meaning or different meanings => Errors, ambiguities and misunderstanding
- Understanding can be specific (“labradoodles have 3 specific fur types”) or general (“mammals produce milk to feed their young”), and aids in bridging the two
  - Situating details and abstractions within a “mental model”
  - Creating a ”whole” (Dictionary.com’s ”synthesis” – ”the combining of the constituent elements … into a single or unified entity; a complex whole formed by combining”)

COMPARING THE SEMANTIC LANDSCAPES

- Both ontologies and LLMs aid in human understanding
  - Ontologies are explicit and structured representations of particular domains (concepts and relationships), with goals of consistency and verifiability
  - Ontologies are valuable for data fusion since terms and concepts from different environments can be mapped to a ”common” ontology
  - LLMs encode linguistic patterns and relationships between tokens/words in the training data, with the goals of producing (reproducing/mimicking) these in its outputs
  - LLMs can analyze images/text/etc. to find similarities and differences, thereby aiding in understanding
- Both can have errors
  - Ontological errors can be introduced by humans when defining the ontology, or when using it (if misapplied due to lack of clarity, misunderstanding of purpose, …)
  - LLM errors can be introduced by inaccuracies and bias in the training data, or by “hallucinations” when outputting results
  - BUT … as black boxes, LLMs are impossible to troubleshoot

LLM AND ONTOLOGY SEMANTIC DIFFERENCES

- Ontologies are conceptual/cognitive/mental/symbolic models of a domain whereas LLMs are statistical models of observed patterns
- Ontologies are designed/curated with specific grounding/reference to their domain, versus LLMs' statistical correlations
  - For example, LLMs recognize "Paris" as statistically associated with the words "city", "France", "Louvre" and "Seine", but have no notion as to the meaning of "city"
- Ontologies model compositionality/relationships (how parts are connected to a whole) whereas LLMs understand wholes/parts in terms of statistical patterns
  - For example, LLMs create readable and logical sentences but may lack coherence and accuracy across those sentences
  - Whereas an ontology indicates that the left rear wheel of a car is "connected to" the left axle that is "connected to" the differential (and so forth), an LLM may state that the wheel is connected to the differential
  - LLMs has no knowledge of how a car is constructed – only of the words that are statistically related to cars and wheels

"KNOWLEDGE" PER CHATGPT

1. **Philosophical Perspective**: Philosophically, knowledge is often defined as justified true belief. This comes from Plato's definition in which, for someone to know something:
   - The person believes it.
   - It is true.
   - The person has good reasons (justification) for believing it.

   However, there are challenges and objections to this definition, such as the Gettier problems, which illustrate situations where one has justified true belief but wouldn't normally say they "know" the proposition.
2. **Information Theory Perspective**: In information theory, knowledge can be seen as organized data or information. Data becomes information when it has meaning, and information becomes knowledge when it is understood in context and is useful.
3. **Practical or Procedural Knowledge**: This is the knowledge of "how" to do something, as opposed to "what" something is. It's more about skills and procedures than facts. Riding a bike is an example of procedural knowledge.
4. **Declarative or Descriptive Knowledge**: This is knowing that something is the case – i.e., knowing facts or "what" something is.
5. **Tacit Knowledge**: This is knowledge that can't be easily articulated or explained. It's the kind of intuitive, hard-to-describe knowledge that might come from personal experience or intuition.
6. **Empirical Knowledge**: This is knowledge based on experience or observation.
7. **Theoretical Knowledge**: Knowledge that is based on theory rather than experience.
8. **Knowledge as Recognition**: Some philosophers argue that knowledge is more about the recognition of patterns or understanding relationships.

In broader terms, knowledge encompasses the facts, information, skills, and understandings that one has acquired through experience and education. It's an awareness or familiarity gained by experience of a fact or situation. It's what is known in a particular field or in total; facts and information stored mentally, as opposed to information that is available in books or other repositories.

Regenerate

COMPARING KNOWLEDGE LANDSCAPES

- Both ontologies and LLMs encode “knowledge”
  - Ontological knowledge is what is chosen and encoded by the designers and implementers
  - LLM knowledge comes from the storage and retrieval of the input data, which can be thought of as representing “human knowledge”
- Neither ontologies nor LLMs possess the ability to ”experience”
  - There is no conscious “understanding” or belief
  - There is no continuous memory, except as may be added by humans and via fine-tuning, retraining, prompt refinement, etc.
- Neither ontologies nor LLMs capture “tacit” knowledge

KNOWLEDGE REPRESENTATION

- "What is Knowledge Representation?" (AI Magazine, Vol 14, 1993)
- Characterized by several roles:
  - Acting as "a *surrogate*, a substitute for the thing itself, … to enable an entity to determine consequences by thinking rather than acting … (by reasoning about the world rather than taking action in it)"
  - With a specific "set of *ontological commitments*", defining what is represented (versus what is ignored)
  - Allowing "*intelligent reasoning*" that is "*efficient*"
    - "… a representation provides for organizing information to facilitate making … inferences"
  - "… a medium of *human expression*, that is, a language in which we say things about the world"

COMPARING REPRESENTATION LANDSCAPES

- Examining the roles (except reasoning/inference which is discussed next):
  - Acting as “a *surrogate*, a substitute for the thing itself, … to enable an entity to determine consequences by thinking rather than acting … (by reasoning about the world rather than taking action in it)”
    - Ontologies relate more closely to surrogates/digital twins than LLMs
  - With a specific “set of *ontological commitments*”, defining what is represented (versus what is ignored)
    - Formally specified ontologies include ontological commitments, although not all “ontologies” have these
    - LLMs have restrictions on what is represented but only by selection/restriction/refinement of training data and prompts
  - “… a medium of *human expression*, that is, a language in which we say things about the world”
    - Ontologies define the ”language” by the concepts that are represented
    - LLMs utilize the patterns in the training data to output ”language”

COMPARING REASONING LANDSCAPES

- Both LLMs and ontologies can be used for reasoning/inference
  - Formal ontologies specify concepts, relationships, axioms and rules that enable reasoning, the detection of inconsistencies and the deduction/inference of new "facts"
  - LLMs *mimic* reasoning and are useful when their training data includes the empirical or procedural knowledge relevant to the prompt
  - "Chain of thought" and "tree of thought" approaches improve LLM results
    - Chain of thought prompts an LLM with explicit steps (or asks the LLM to define these steps) for logical coherence, and then executes them in order
    - Tree of thought prompts an LLM to explore multiple reasoning paths and then evaluate the likelihoods/pros/cons/… of each
- Both LLMs and ontologies may have reasoning errors
  - Ontologies cannot be used for reasoning if they contain inconsistencies, or have limited reasoning capabilities due to inadequate or incomplete specifications
  - Reasoning over complex ontologies can be computationally expensive
  - LLMs do not use formal logic or a reasoning process, and results can be sensitive to small variations/changes to the input prompts as well as the possibility of producing different answers even if the same prompts are reused

LLM REASONING RESEARCH

- Large Language Models are In-Context Semantic Reasoners rather than Symbolic Reasoners (June 2023)
  - "… emergent ability of LLMs to utilize a "zero-shot" or "few-shot" learning approach without any gradient updates—a task description or a few examples are provided to guide their reasoning process. One typical example is the "chain-of-thought (CoT)" approach …"
  - Findings:
    - "When semantics are consistent with commonsense, LLMs perform fairly well; when semantics are decoupled or counter-commonsense, LLMs struggle to solve the reasoning tasks by leveraging in-context new knowledge."
  - Recommendations:
    - "More complex symbolic reasoning benchmark: … developing new datasets with decoupled semantics and more complex reasoning tasks … These benchmarks should challenge LLMs with diverse and intricate symbolic knowledge."
    - "Combination with external non-parametric knowledge base: … the memorization abilities of LLMs are not comparable to existing graph-based methods. This motivates integrating LLMs with external non-parametric knowledge bases, such as graph databases, to enhance their knowledge insertion and updating. This **hybrid approach** can leverage the strengths of LLMs' language understanding and the comprehensive, accurate and up-to-date knowledge stored in non-parametric sources."
    - "Improving the ability of processing in-context knowledge: … improve LLMs' capabilities in processing and leveraging in-context knowledge … includes developing mechanisms to better encode and retrieve relevant information"

HYBRID SYSTEMS – COMBINING THE TECHNOLOGIES

- Examples:
  - LLMs can be used to extract information from text and aid in its mapping to an ontology/in creating instances in a knowledge graph
  - LLMs can make ontologies more accessible and usable for non-expert users - For example, by generating SPARQL queries from natural language or by summarizing a knowledge graph
  - LLMs can aid in searching for related concepts defined in other ontologies and knowledge graphs
  - Ontologies can be used to formulate prompts to an LLM, or to validate the responses of an LLM

HYBRID SYSTEMS – COMBINING THE TECHNOLOGIES

- Some thoughts from Tony Seale, https://www.linkedin.com/feed/update/urn:li:activity:7088065122153177089/
  - "…when deploying LLMs within an enterprise context, reliability, trustworthiness, and understandability are vital concerns for those running and governing these systems. Hallucination is simply not an option."
  - "The LLMs help discover new knowledge, and the ontologies compile that knowledge down for future use"
  - "As LLMs help generate better ontologies faster and more dynamically, the ontologies, in turn, elevate the performance of LLMs by offering a more comprehensive context of the data and text they analyse."

HYBRID SYSTEMS – COMBINING THE TECHNOLOGIES

- Some thoughts from John Sowa (email):
  - "… not saying that LLMs are incompatible with ontologies. But I am emphasizing that anything based on natural languages (as expressed in LLMs) will be limited to the level of precision expressed in typical NL texts. Nobody would fly in an airplane whose specifications were limited to the kinds of NL text that describes an airplane."
  - "That is why we need formal ontologies. They can be as detailed and precise as required, and they can be tested by the most fundamental computational methods available today or imaginable in the future."
  - "LLMs are based on trillions of bytes of vague, inconsistent, and unreliable natural language text. Some of it may be good. Some of it may be brilliant. But none of it can match the level of precision required for designing an airplane, a rocket to the moon, or an operation on a human heart."
  - "… most general-purpose ontologies don't go beyond the precision that can be expressed in ordinary NL text. But special-purpose ontologies can be made as precise and formal as required for the applications."

SESSION AND SPEAKER OVERVIEW

- 11 October 2023, *Setting the stage*
  - **Deborah McGuinness**, Rennselaer Tetherless World Senior Constellation Chair, Professor of Computer Science, Cognitive Science, and Industrial and Systems Engineering
  - The Evolving Landscape: Generative AI, Ontologies, and Knowledge Graphs
- 18 and 25 October 2023, *A look across the industry*
  - **Kurt Cagle**, Author of the Cagle Report
  - **Tony Seale**, Knowledge graph architect and thought leader (LinkedIn)
  - **Evren Sirin**, Stardog CTO and lead of the Voicebox offering
  - **Yuan He**, Key contributor to DeepOnto, a package for ontology engineering with deep learning
- 1 November 2023, *Demos of hybrid systems*
  - **Andrea Westerinen**, Deep Narrative Analysis
  - **Prasad Yalamanchi**, Lead Semantics CTO
- 8 November 2023, *Broader thoughts*
  - **Anatoly Levenchuk**, *Hybrid reasoning, the scope of knowledge, and what is beyond ontologies?*
  - **John Sowa** and **Arun Majumdar**, *LLMs, ontologies, and formal systems*
- 15 November 2023, *Discussion and Synthesis*