Neurosymbolic Graph Enrichment for Grounded World Models

Stefano De Giorgis[1], Aldo Gangemi[2,3], and Alessandro Russo[2]

[1] Vrije Universiteit, Amsterdam, Netherlands
s.degiorgis@vu.nl
[2] Institute for Cognitive Sciences and Technologies - National Research Council (CNR-ISTC), Italy
[3] Department of Philosophy - University of Bologna, Italy

The work, published on IPM - Information Processing and Management journal[4], presents Polanyi, a novel neuro-symbolic framework that addresses a current challenge in artificial intelligence: developing systems capable of reasoning about complex real-world scenarios with human-like comprehension. While Large Language Models (LLMs) excel at pattern recognition and generation, and knowledge-based systems provide structured reasoning capabilities, neither approach alone sufficiently captures the multifaceted nature of human understanding required for Grounded World Models (GWMs) [2, 5] . Our approach bridges this gap by leveraging LLMs not as expert systems but as reactive engines to extract implicit contextual commonsense knowledge.

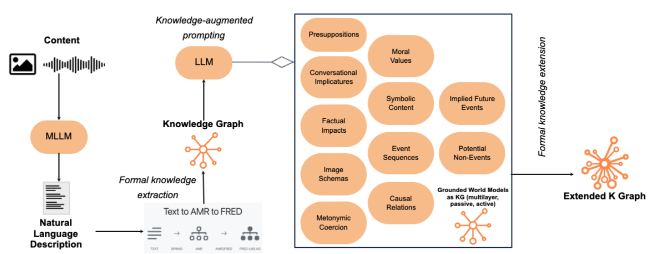

Fig. 1. Polanyi's hybrid knowledge enrichment pipeline.

The framework implements a modular pipeline that transforms multimodal inputs (text or images) into enriched knowledge graphs through several key stages: (1) multimodal LLMs generate natural language descriptions from images, (2) the Text2AMR2FRED tool [3] converts descriptions into Abstract Meaning Representation (AMR) graphs and subsequently into OWL-RDF knowledge graphs with alignments to public resources like WordNet [8], PropBank [6], and DOLCE [1], and (3) LLMs iteratively extend these base graphs with implicit knowledge across 11 distinct heuristics. The 11 heuristics capture diverse aspects of tacit knowledge essential for human-like understanding. The list

[4] https://www.sciencedirect.com/science/article/pii/S030645732500069X

includes: *Presuppositions*, based on previous background knowledge; *Conversational Implicatures* [4], which often contributes in making sense of incomplete information in linguistic exchanges; *Factual Impact*, which grounds linguistic entities to factual knowledge; *Image Schemas*, basic building blocks of cognition which grounds our way of conceiving the world in our sensori-motor bodily perception [7]; *Metonymic Coercions*, which allows understand propositions whose truth value would be zero, but differ from metaphorical speech grounding for the parthood-whole relation; *Moral Value Driven Coercion*, applied everyday in appraisal and moral evaluative processes, values nudge our daily behavior; *Symbolic Coercion*, in Peirce terminology [9], used to anchor meaning to various entities of the world; *Event Sequences*, determinant in our plan-making capability and ability to design plausible scenarios and outcomes; *Causal Relations*, establishing relations of cause-effect between processes and events, to avoid either having only (i) temporal sequences and (ii) statistical correlation; *Implied Future Events*, a specification of Event Sequences, for temporal projection in the future; and *Implied Non-Events*, an infinite set of events, but, referring to the Frame problem, focusing on those more closely related to a specific Event Sequence. We conducted comprehensive evaluation across three experiments using rigorous methodologies. Experiment 1 demonstrated the complete pipeline on sports imagery, generating a base graph, and extending it with LLM enrichment iterations, ranging from 12 to 63 triples per heuristic. Human evaluation using 5-point Likert scales revealed high plausibility mean ratings ($\mu > 3.0$) for all heuristics, with Factual Impact, Conversational Implicatures, and Moral Value-driven Coercions achieving particularly strong performance ($\mu > 4.29$). Logical validation using HermiT reasoner confirmed structural integrity across all generated graphs. Experiment 2 adopted LLMs as judges, and compared knowledge generation capabilities across three state-of-the-art models (Claude 3.5 Sonnet, GPT-4o, Mistral Large 2) using the same base scenario. Results revealed significant productivity variations, with Mistral generating over 100 triples for certain heuristics while maintaining consistent quality. Self-coherence analysis showed Claude and Mistral achieving robust consistency ($>$0.8), while inter-model agreement remained moderate (Krippendorff's $\alpha < 0.5$), indicating distinct but valid evaluation criteria across architectures. Finally, Experiment 3 validated practical applicability through a downstream task predicting plausible future events from 12 recent New York Times articles. Our system achieved 100% precision in capturing LLM predictions while providing a Structure Multiplication Factor of 2.74, demonstrating superior semantic granularity compared to natural language predictions. The framework's modular architecture offers significant advantages over monolithic approaches: individual components can be updated independently, enabling integration of advances in LLMs, AMR parsing, and entity linking without system-wide modifications. This design ensures adaptability to diverse domains while maintaining transparency and interpretability through inspectable intermediate outputs.

References

1. Stefano Borgo, Roberta Ferrario, Aldo Gangemi, Nicola Guarino, Claudio Masolo, Daniele Porello, Emilio M Sanfilippo, and Laure Vieu. Dolce: A descriptive ontology for linguistic and cognitive engineering. *Applied ontology*, 17(1):45–69, 2022.
2. Jay W Forrester. Counterintuitive behavior of social systems. *Theory and decision*, 2(2):109–140, 1971.
3. Aldo Gangemi, Arianna Graciotti, Antonello Meloni, Andrea Giovanni Nuzzolese, Valentina Presutti, Diego Reforgiato Recupero, Alessandro Russo, and Rocco Tripodi. Text2AMR2FRED, a Tool for Transforming Text into RDF/OWL Knowledge Graphs via Abstract Meaning Representation. In *Proceedings of the ISWC 2023 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 22nd International Semantic Web Conference (ISWC 2023)*, 2023.
4. HP Grice. Logic and conversation. *Syntax and semantics*, 3, 1975.
5. David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018.
6. Paul R Kingsbury and Martha Palmer. From treebank to propbank. In *LREC*, pages 1989–1993, 2002.
7. George Lakoff, Mark Johnson, et al. *Philosophy in the flesh: The embodied mind and its challenge to western thought*, volume 640. Basic books New York, 1999.
8. George A Miller. *WordNet: An electronic lexical database.* MIT press, 1998.
9. Charles Sanders Peirce and Justus Buchler. Logic as semiotic: The theory of signs. *Philosophical Writings of Peirce, ed. Justus Buchler (New York: Dover, 1955)*, page 100, 1902.