

Prompt-Time Ontology-Driven Symbolic Knowledge Capture with Large Language Models

Tolga Çöplü*, Arto Bendiken, Andrii Skomorokhov, Eduard Bateiko and Stephen Cobb

Haltia, Inc.

Abstract

In applications such as personal assistants, large language models (LLMs) must consider the user's personal information and preferences. However, LLMs lack the inherent ability to learn from user interactions. This paper explores capturing personal information from user prompts using ontology and knowledge-graph approaches. We use a subset of the KNOW ontology, which models personal information, to train the language model on these concepts. We then evaluate the success of knowledge capture using a specially constructed dataset. Our code and datasets are publicly available at <https://github.com/HaltiaAI/paper-PTODSKC>

Keywords

ontology-driven symbolic knowledge capture, KNOW ontology, symbolic representation, knowledge graphs, large language models, fine-tuning

1. Introduction

Currently, many generative artificial intelligence (AI) applications, particularly personal assistants, strive to offer users personalized experiences. To achieve this, AI applications must learn personal information and preferences from user interactions (knowledge capture) and use this learned knowledge in future conversations (knowledge utilization). Implementing this fundamental personal AI approach depends on addressing several complex sub-problems, such as discerning which user prompt information is personal, extracting it, determining whether the extracted information is duplicate, and associating it with other personal data.

These challenges have been the focus of extensive research within the AI field for many years. However, the emergence of neurosymbolic approaches through the collaboration between large language models (LLMs) and symbolic AI has provided researchers with new perspectives [1, 2, 3, 4]. LLMs' capabilities in natural language processing can be integrated with the representational and factual reasoning abilities of knowledge graphs, enhanced by the structure, rules, and inference mechanisms offered by an ontology. For targeted personal AI applications, this ontology approach presents several benefits:

- Ontology schemas enable language models to de-

termine which personal information will be captured and how it will be associated with other captured knowledge.

- Ontology rules can help identify inconsistencies in the captured knowledge, allowing for validation before storage.
- Ontology relationships allow the extraction of implicit information from captured knowledge, effectively enabling automatic inference that expands the knowledge graph.
- A robust, personalized knowledge graph forms a reliable foundation for facilitating personalized interactions with the application through language models.

In this paper, we address a specific aspect of the AI personalization challenge by focusing on prompt-time, ontology-driven symbolic knowledge capture using language models. We explore the extraction from user prompts of subject-predicate-object triples¹ that conform to a specified ontology. We have investigated various methods to enable the underlying language model to comprehend a pre-defined ontology, ensuring effective symbolic knowledge capture. By utilizing a specially designed dataset, we evaluate the effectiveness of these methods, emphasizing their strengths and identifying potential areas for improvement.

The structure of this paper is as follows: Section 2 discusses in-context learning and fine-tuning approaches for ontology-driven symbolic knowledge capture and focuses on the details of the fine-tuning approach. Section 3 describes the experimental setup by presenting the development framework, the language model selection, and the ontology and dataset creation process. Section 4

KiL'24: Workshop on Knowledge-infused Learning co-located with 30th ACM KDD Conference, August 26, 2024, Barcelona, Spain

* Corresponding author.

✉ tolga@haltia.ai (T. Çöplü); arto@haltia.ai (A. Bendiken); andriy@haltia.ai (A. Skomorokhov); eduard@haltia.ai (E. Bateiko); steve@haltia.ai (S. Cobb)

🆔 0009-0004-9414-0588 (T. Çöplü); 0009-0002-0725-4874 (A. Bendiken); 0000-0002-5696-6723 (A. Skomorokhov); 0000-0002-6729-5611 (E. Bateiko); 0009-0004-0476-6000 (S. Cobb)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://www.w3.org/TR/rdf12-concepts/>

outlines our performance evaluation framework and the test results. Finally, Section 5 concludes the paper and suggests future directions.

2. Ontology-Driven Symbolic Knowledge Capture

In the literature, language models have demonstrated their capability to transform unstructured text into a knowledge graph [5, 6, 7, 8, 9]. However, the process of populating a knowledge graph from user prompts in alignment with a pre-defined ontology has been explored only marginally [10, 11, 12, 13, 14]. Research typically centers on in-context learning, which heavily relies on prompt engineering. A significant limitation of this approach is the requirement to incorporate the entire custom ontology into the prompt. This necessity not only slows down the knowledge capture process, because of the high token overhead but also restricts the use of larger ontologies due to the constraint on context-window length. Given these constraints, in-context learning methods do not provide a scalable solution for ontology-driven symbolic knowledge capture.

An alternative approach involves training a language model with a pre-defined ontology, so that the model internalizes it. There are two strategies to consider: pre-training the LLM on the ontology or fine-tuning it. This paper does not explore pre-training due to its extensive data, computational resources, energy, and time requirements. Additionally, pre-training does not offer a flexible response to ongoing changes or expansions in the ontology. Therefore, this paper will focus on fine-tuning as a method to train language models on personal ontologies, highlighting advantages in feasibility and maintainability.

2.1. Ontology-Driven Knowledge Capture with Fine-Tuning

Fine-tuning is a process whereby a pre-trained language model is further trained on a specific dataset to tailor its capabilities to a particular task. In our study, the language model is expected to learn the classes, object properties, and data properties defined in an ontology, and to use them to populate a knowledge graph from user prompts. The first step involves preparing a fine-tuning dataset, which includes user prompts, system prompts, and expected model responses for each concept in the ontology. This dataset is used to fine-tune the language model, which is then evaluated by testing it with new prompts to assess the effectiveness of the knowledge capture operation. We define a system prompt for this task with the requirement of maintaining the model’s generality across other tasks.

The following points highlight the key aspects of ontology fine-tuning:

- The training dataset’s coverage and diversity are vital for successful fine-tuning. These characteristics greatly influence the LLM’s ability to effectively capture knowledge. Details about the dataset and how it is constructed are discussed in Section 3.4.
- The training dataset must include a variety of examples for each element of the predefined ontology. This approach avoids scalability issues typically associated with in-context learning and ensures comprehensive learning coverage.
- If the LLM encounters a user prompt that is not relevant to the predefined ontology concepts, it should not attempt to capture knowledge. Therefore, the dataset should also contain sufficient out-of-context samples to enable the LLM to distinguish between relevant and irrelevant information for capture.

3. Experimental Setup

This section explores the components of our experimental setup.

3.1. Development Framework

The methods suggested in this paper have been implemented using the Apple MLX framework [15]. MLX is a specialized array framework designed for machine learning applications, akin to NumPy, PyTorch, or JAX, with the distinction of being exclusive to Apple silicon.

Ontology fine-tuning has been conducted using the parameter-efficient QLoRA approach [16] on our custom dataset, comprising randomly selected, non-overlapping sets of training, validation, and test samples.

3.2. Language Model

The methods we have developed here do not have a structural dependency on a particular underlying foundation model. The key factors guiding our language model selection were its proven effectiveness across diverse domains in community benchmarks and its prevalence in the field. Owing to its performance in the Hugging Face Open LLM Leaderboard [17] and its robust ecosystem, the Mistral-7B-Instruct-v0.2 [18], based on the Llama 2 [19] architecture, was selected for our research. We ran all examples, tests, and benchmarks on the MLX 4-bit quantized version of this model to be able to run the tests on personal laptops.

3.3. Predefined Ontology

Our study is inspired by KNOW[20]—the Knowledge Navigator Ontology for the World—and utilizes it for representing personal information. KNOW is introduced as a pioneering framework designed to capture everyday knowledge to enhance language models in real-world generative AI applications such as personal AI assistants. The ontology focuses on human life, encompassing everyday concerns and significant milestones, and limits its initial scope to established human universals, including spacetime (places, events) and social dimensions (people, groups, organizations). This pragmatic approach emphasizes universality and utility, contrasting with previous works like Schema.org[21] and Cyc[22] by building on language models’ inherent encoding of salient common-sense knowledge.

Due to the requirement that each element in the ontology be associated with a diverse set of prompt and response samples within the training dataset, our research focuses on a specific subset of the KNOW ontology. This subset concentrates on core family relationships with four ontology classes, eleven object properties, and one data property. A visual depiction of this subset is presented in Figure 1.

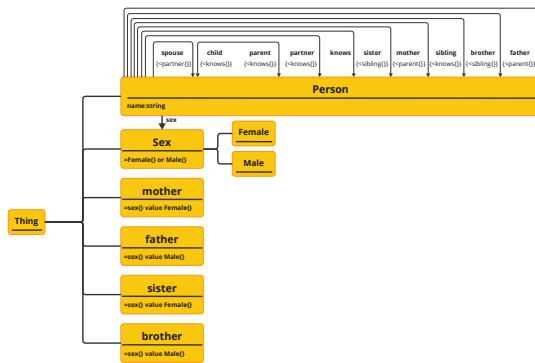


Figure 1: A visual representation of the ontology design used in this paper.

3.4. Dataset

For a language model to effectively learn a predefined ontology and use it to perform knowledge extraction and capture, a robust and diverse training dataset is essential. Our paper focuses on a subset of the KNOW ontology that includes the concepts of ‘person’, ‘name’, ‘sex’, ‘child’, ‘father’, ‘mother’, ‘sibling’, ‘sister’, ‘brother’, ‘spouse’, ‘partner’ and ‘knows’. We created 143 manually crafted user prompts along with their respective ontology responses for training and tests. Additionally, to manage inputs that fall outside these ontology concepts,

we included 32 generic user prompts in the dataset. The composition of this dataset, which consists of 175 user prompts, is illustrated in Figure 2. Concepts not associated with the ontology are labeled as the ‘none’ legend in the figure. As each sample prompt typically contains multiple modeled concepts, the chart shows a total number of concept occurrences greater than the number of prompts.




Figure 2: Occurrences of ontology concepts in the prepared dataset.

The Turtle format was chosen for serializing the ontology population in our research because of its straightforward structure, readability, and prevalent use in existing pre-training datasets for LLMs.

4. Performance Evaluation

Our research focuses on fine-tuning a language model with predefined ontology concepts and capturing knowledge from user prompts that fits the ontology. This section will detail the performance evaluations associated with these efforts.




Figure 3: Occurrences of ontology concepts in the test dataset.

Initially, we investigated how many training samples each ontology concept required to effectively teach the

core family ontology to the selected Mistral-7B-Instruct-v0.2 model. From the dataset described in Section 3.4, 41 random samples were selected and reserved for method evaluation. The distribution of ontology concepts within this test set is shown in Figure 3.

From the remaining 134 samples, we created three different training datasets. In these three datasets, we have ensured the inclusion of 2, 4, and 8 sample prompts for the ontology concepts 'child', 'father', 'mother', 'sibling', 'sister', 'brother', 'spouse', 'partner', and 'knows'. Subsequent evaluations using the test set measured the precision, recall, and f1-scores for each fine-tuning session. During these evaluations, the generated prompt responses were processed triple by triple and compared against the ground truth established for the test set. The findings are displayed in Figure 4.




Figure 4: Ontology population performance (precision, recall, and f1-score) for various fine-tuning dataset sizes.

The tests were conducted using the default QLoRA hyperparameters specified in the MLX framework. To ensure consistent test results, each training session was set to run for 18 epochs.

- Layer keys to apply: ["self_attn.q_proj", "self_attn.v_proj"]
- Rank: 8
- Alpha: 16
- Scale: 10
- Optimizer: Adam
- Learning rate: 1×10^{-5}
- Number of layers to fine-tune: 16
- Minibatch size: 4

As illustrated in Figure 4, our study, which encompasses twelve key ontology concepts, demonstrates that providing eight diverse examples for each concept yields acceptable success rates. Although our training and test datasets are not sufficiently diverse or large enough to generalize the results to real user scenarios, the high success achieved with a small number of training samples is promising for the feasibility of the proposed approach.

In the subsequent phase, we explored the optimal number of training epochs required to achieve maximum performance for the training set. For this analysis, we continued using the default MLX QLoRA hyperparameters with the 8 samples per ontology concept, but trained the QLoRA adapter over various epoch lengths. We then conducted evaluations on the test set using each trained adapter, and the findings are presented in Figure 5.




Figure 5: Ontology population performance (precision, recall, and f1-score) of 8 samples per ontology concept for various epochs.

As depicted in Figure 5, the success rate of the ontology population increases with longer training. However, considering the resource usage and energy consumption, we observe that 18 epochs are sufficient for fine-tuning.

5. Conclusion

In this paper, we first explored prompt-driven, ontology-based symbolic knowledge capture and its importance in the generative AI domain. We then discussed the ontology approach and how to teach ontology concepts to the language model through in-context learning and training. The language model was fine-tuned using a custom dataset focused on core family relationships, and we evaluated the model's ability to learn personal ontology concepts.

Our findings indicate that fine-tuning is particularly effective for teaching ontology concepts to a language model for prompt-time knowledge capture. In our future work, we aim to integrate the generated knowledge graph with the language model for knowledge utilization, combining the strengths of the neural and symbolic AI approaches.

Please refer to the paper's corresponding GitHub repository at <https://github.com/HaltiaAI/paper-PTODSKC>

References

- [1] A. Sheth, K. Roy, M. Gaur, Neurosymbolic AI – Why, What, and How, 2023. URL: <http://arxiv.org/abs/2305.00813>, arXiv:2305.00813 [cs].
- [2] L. N. DeLong, R. F. Mir, J. D. Fleuriot, Neurosymbolic AI for Reasoning over Knowledge Graphs: A Survey, 2024. URL: <http://arxiv.org/abs/2302.07200>, arXiv:2302.07200 [cs, stat].
- [3] F. J. Ekaputra, M. Llugiqi, M. Sabou, A. Ekelhart, H. Paulheim, A. Breit, A. Revenko, L. Waltersdorfer, K. E. Farfar, S. Auer, Describing and Organizing Semantic Web and Machine Learning Systems in the SWeMLS-KG, 2023. URL: <http://arxiv.org/abs/2303.15113>, arXiv:2303.15113 [cs].
- [4] L.-P. Meyer, C. Stadler, J. Frey, N. Radtke, K. Jung-hanns, R. Meissner, G. Dziwis, K. Bulert, M. Martin, LLM-assisted Knowledge Graph Engineering: Experiments with ChatGPT, 2023. URL: <http://arxiv.org/abs/2307.06917>, arXiv:2307.06917 [cs].
- [5] Q. Guo, Z. Jin, X. Qiu, W. Zhang, D. Wipf, Z. Zhang, CycleGT: Unsupervised Graph-to-Text and Text-to-Graph Generation via Cycle Training, 2020. URL: <http://arxiv.org/abs/2006.04702>. doi:10.48550/arXiv.2006.04702, arXiv:2006.04702 [cs].
- [6] X. Li, F. Yin, Z. Sun, X. Li, A. Yuan, D. Chai, M. Zhou, J. Li, Entity-Relation Extraction as Multi-Turn Question Answering, 2019. URL: <http://arxiv.org/abs/1905.05529>. doi:10.48550/arXiv.1905.05529, arXiv:1905.05529 [cs].
- [7] Y. Xu, L. Fu, Z. Lin, J. Qi, X. Wang, INFINITY: A Simple Yet Effective Unsupervised Framework for Graph-Text Mutual Conversion, 2022. URL: <http://arxiv.org/abs/2209.10754>. doi:10.48550/arXiv.2209.10754, arXiv:2209.10754 [cs].
- [8] R. Anantharangachar, S. Ramani, S. Rajagopalan, Ontology Guided Information Extraction from Unstructured Text, *International journal of Web & Semantic Technology* 4 (2013) 19–36. URL: <http://arxiv.org/abs/1302.1335>. doi:10.5121/ijwest.2013.4102, arXiv:1302.1335 [cs].
- [9] T. Çöplü, A. Bendiken, A. Skomorokhov, E. Bateiko, S. Cobb, J. J. Bouw, Prompt-Time Symbolic Knowledge Capture with Large Language Models, 2024. URL: <http://arxiv.org/abs/2402.00414>. doi:10.48550/arXiv.2402.00414, arXiv:2402.00414 [cs].
- [10] H. B. Giglou, J. D’Souza, S. Auer, LLMs4OL: Large Language Models for Ontology Learning, 2023. URL: <http://arxiv.org/abs/2307.16648>. doi:10.48550/arXiv.2307.16648, arXiv:2307.16648 [cs, math].
- [11] P. Mateiu, A. Groza, Ontology engineering with Large Language Models, 2023. URL: <http://arxiv.org/abs/2307.16699>, arXiv:2307.16699 [cs].
- [12] K. Lu, X. Pan, K. Song, H. Zhang, D. Yu, J. Chen, PIVOINE: Instruction Tuning for Open-world Information Extraction, 2023. URL: <http://arxiv.org/abs/2305.14898>. doi:10.48550/arXiv.2305.14898, arXiv:2305.14898 [cs].
- [13] N. Mihindukulasooriya, S. Tiwari, C. F. Enguix, K. Lata, Text2KGBench: A Benchmark for Ontology-Driven Knowledge Graph Generation from Text, 2023. URL: <http://arxiv.org/abs/2308.02357>. doi:10.48550/arXiv.2308.02357, arXiv:2308.02357 [cs].
- [14] M. Funk, S. Hosemann, J. C. Jung, C. Lutz, Towards Ontology Construction with Language Models, 2023. URL: <http://arxiv.org/abs/2309.09898>. doi:10.48550/arXiv.2309.09898, arXiv:2309.09898 [cs].
- [15] A. Hannun, J. Digani, A. Katharopoulos, R. Collobert, MLX: Efficient and flexible machine learning on Apple silicon, <https://github.com/ml-explore>, 2023.
- [16] T. Dettmers, A. Pagnoni, A. Holtzman, L. Zettlemoyer, QLoRA: Efficient Finetuning of Quantized LLMs, 2023. URL: <http://arxiv.org/abs/2305.14314>. doi:10.48550/arXiv.2305.14314, arXiv:2305.14314 [cs].
- [17] E. Beeching, C. Fourrier, N. Habib, S. Han, N. Lambert, N. Rajani, O. Sanseviero, L. Tunstall, T. Wolf, Open LLM Leaderboard, https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard, 2023.
- [18] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. I. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mistral 7B, 2023. URL: <http://arxiv.org/abs/2310.06825>. doi:10.48550/arXiv.2310.06825, arXiv:2310.06825 [cs].
- [19] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, T. Scialom, Llama 2: Open foundation and fine-tuned chat models, 2023. arXiv:2307.09288.

- [20] A. Bendiken, KNOW: A Real-World Ontology for Knowledge Capture with Large Language Models, 2024. URL: <https://arxiv.org/abs/2405.19877>. arXiv:2405.19877.
- [21] R. V. Guha, D. Brickley, S. Macbeth, Schema.org: evolution of structured data on the web, *Communications of the ACM* 59 (2016) 44–51. URL: <https://doi.org/10.1145/2844544>. doi:10.1145/2844544.
- [22] D. B. Lenat, R. V. Guha, *Building Large Knowledge-Based Systems; Representation and Inference in the Cyc Project*, 1st ed., Addison-Wesley Longman Publishing Co., Inc., USA, 1989.