

Knowledge Graph Completion for Action Prediction on Situational Graphs

A Case Study on Household Tasks

Mariam ARUSTASHVILI^{a,1}, Jörg DEIGMÖLLER^b and Heiko PAULHEIM^c

^a*University of Stavanger, Information Access and Artificial Intelligence group, Norway*

^b*Honday Research Institute Europe, Germany*

^c*University of Mannheim, Data and Web Science Group, Germany*

Abstract. Knowledge Graphs are used for various purposes, including business applications, biomedical analyses, or digital twins in industry 4.0. In this paper, we investigate knowledge graphs describing household actions, which are beneficial for controlling household robots and analyzing video footage. In the latter case, the information extracted from videos is notoriously incomplete, and completing the knowledge graph for enhancing the situational picture is essential. In this paper, we show that, while a standard link prediction problem, situational knowledge graphs have special characteristics that render many link prediction algorithms not fit for the job, and unable to outperform even simple baselines.

Keywords. Knowledge Graph, Completion, Link Prediction, Situational Knowledge Graph, Household Actions, Robotics

1. Introduction

Knowledge graphs (KGs) are structured representations of knowledge that model entities and their relationships in a graph format, offering a semantically rich and machine-readable alternative to traditional databases [10,12]. By capturing complex relational data through nodes and edges, KGs enable advanced reasoning, inference, and data integration across heterogeneous sources [26]. Their versatility has made them essential in fields such as natural language processing, biomedical research, and personalized recommendation systems [13,25,40]. Specialized forms like situational and temporal knowledge graphs extend this functionality by incorporating dynamic or context-dependent information, supporting applications in rapidly changing environments such as robotics and context-aware systems [5,15,17,43].

In robotics, KGs provide a structured framework for representing and reasoning about environments, enabling robots to interpret human actions, plan tasks, and adapt to new situations [2]. Projects like RoboEarth [39] and KnowRob [36] show how perception and action knowledge can be integrated to support autonomous behavior. Ontology-based approaches further highlight the value of structured reasoning for robotic percep-

¹Corresponding Author: Mariam Arustashvili; University of Stavanger; E-mail: mariam.arustashvili@uis.no.

tion, planning, and decision-making [27]. Recent advances also show how multi-modal knowledge integration enhances embodied AI [34].

A central challenge in robotics is reasoning under uncertainty. Robots operating in real-world environments—such as homes or workplaces—often face incomplete data due to occlusions, sensor limitations, or dynamic change [29]. Knowledge Graph Completion (KGC) addresses this by inferring missing links (e.g., unobserved object interactions) or entities (e.g., likely tools) within structured graphs [30,33].

Another challenge is learning patterns in action sequences. Robots must predict the next steps of human behavior based on context, rather than relying solely on common-sense assumptions or static rules. In this work, we investigate how well KGC methods support this task by testing them on situational KGs built from household activity videos. Specifically, we evaluate how different models predict a human’s goal (parent action) or next step (sub-action) in a sequence—e.g., predicting milk handling after observing cereal pouring—so the robot can assist appropriately.

We compare standard link prediction (LP) models (e.g., TransE, ComplEx), simple statistical baselines, and large language models (LLMs). Results show that conventional embedding models struggle to outperform even simple heuristics, likely due to the ambiguity and temporal nature of human activities. While LP remains central to KGC, our findings suggest that LLMs—thanks to their contextual reasoning and generalization abilities—may offer a more robust solution for parent action tasks. In contrast, sub-action prediction presents a greater challenge for LLMs and LP models, where simpler, pattern-based statistical methods outperform more complex approaches. We conclude with a discussion of open challenges and the potential of hybrid approaches.

2. Related Work

Using knowledge graphs to represent information about the surroundings is common in robotics [32]. Semantic knowledge is essential for robots autonomy [27], scene understanding in static and dynamic environments, and during human-machine interaction [11]. Therefore, since graph-based data is used for knowledge representation in robots, it is common to make inferences such as action prediction on those situational graphs.

A prominent example of structured knowledge representation in robotics is the Functional Object-Oriented Network (FOON) [31], which encodes task-related object and action knowledge using a bipartite graph structure. While our graph shares some structural similarities with FOON in terms of capturing object-action interactions, it differs in key ways: we emphasize temporal and hierarchical relations between sub-actions (e.g., "has_next" and "has_element") rather than focusing on functional units of manipulation. Additionally, our approach is data-driven, relying on per-frame annotations from the KIT Bimanual Actions Dataset, whereas FOON is often constructed manually or semi-manually.

Datasets that capture human activities with rich temporal and object-level detail are essential for advancing KGC and action prediction in robotics. Not many public datasets are available that allow experimentation with the action inference task using link prediction methods. Table 1 shows the requirements that guided our dataset choice. The KIT Bimanual Actions Dataset [8] is a prominent example, originally developed for computer vision tasks and widely used in research on human-object interaction and bimanual ac-

Dataset	Parent actions	Sequential sub-actions	Time-stamps	Objects	States/ Properties
The Breakfast Actions Dataset [19] [20]	Yes	Yes	Yes	Not directly, but possible to extract	No
KIT Bimanual Actions Dataset [8]	Yes	Yes	Yes	Yes	Yes
AVA-Kinetics Dataset [21]	Yes	Yes	Yes	No	No
EPIC-KITCHENS-100 [6]	No	Yes	Yes	Yes	No
Georgia Tech Egocentric Activity Datasets [22]	Yes	Yes	Yes	Yes	No
YouCook [7]	Yes	Yes	Yes	Yes	No

Table 1. Comparison of various action recognition datasets based on their features, including parent actions, sequential sub-actions, timestamps, objects, and states/properties.

tion recognition. It consists of 540 RGB-D recordings across nine tasks performed by six subjects, with detailed, frame-level annotations of sub-actions for both hands. These annotations include spatial object data and allow for modeling fine-grained action sequences relevant to robotic planning and prediction.

[42] investigates predicting subsequent activities and estimating the time until an activity requiring assistance is initiated. This task is similar to what we do in this paper; however, they use a different methodology, namely Markov chains, to model human behavior and test it in a real-world setting. Literature is sparse on human action prediction using link prediction techniques. [38] discusses signaling in human-robot interaction as a tool to predict human actions. Other researchers perform activity prediction directly on visual data from videos [16,24]. In contrast, our work models the data derived from video input as a structured knowledge graph and explores link prediction methods on two tasks: (1) subsequent action prediction and (2) parent action prediction, i.e., predicting the overall task of an observed sequence of activities.

3. Experiments

This study leverages the KIT Bimanual Actions Dataset [8], a structured benchmark originally designed for computer vision but well-suited for robotic action prediction due to its granular labeling of bimanual tasks. The dataset contains 540 recordings in which six subjects perform nine distinct tasks (parent actions)—such as preparing cereal or assembling tools—carried out in kitchen and workshop scenarios. Each task is repeated ten times and decomposed into sequential, fine-grained sub-actions (e.g., "approach bowl," "stir with whisk"). It provides frame-level ground truth for these sub-actions alongside 3D object interactions, enabling temporal and spatial modeling of action sequences. Key data is organized hierarchically in JSON files, with per-hand sub-action sequences, including start and end timestamps, and 3D bounding boxes for both hands and objects. During preprocessing, object-sub-action associations were established via bounding box

For consistent and fair evaluation, we applied a fixed train–test split: two out of ten repetitions per task were reserved for testing. This yields 108 graph components in the test set, with 11,088 nodes and 25,897 edges in training, and 2,840 nodes and 6,680 edges in testing. This split is held constant across all models. The processed dataset, splits, and code are publicly available via the MemNet repository.²

We implemented baseline models for both parent action and sub-action prediction tasks to serve as performance benchmarks. These baselines are based on simple statistical assumptions. In both tasks, Baseline1 uses only action information, while Baseline2 incorporates contextual object information, providing a more nuanced view of the activity.

For the parent action prediction task, Baseline1 maps each sub-action to its most frequently co-occurring parent action in the training data. For each activity (i.e., a sequence of sub-actions), the final parent action is determined via majority vote across predicted labels. Baseline2 extends this by using pairs of sub-actions and their associated objects (e.g., "cut–knife", "hold–whisk") to build a more detailed mapping. Including object context significantly improves performance, as shown in Table 3.

For the subsequent action prediction task, Baseline1 predicts the most probable next sub-action based on the current one, using the most frequent transitions observed in training. Baseline2 considers the combination of the current sub-action and its associated object when determining the next step.

For knowledge graph completion, we evaluated several widely-used embedding-based link prediction models using the PyKeen framework [1]: TransE [3], TransR [23], ComplEx [37], DistMult [41], and RotatE [35]. Since action and sub-action entities are literal strings, we also tested LiteralE variants of DistMult and ComplEx[18], which incorporate textual literals into entity representations.

In addition to traditional KG models, we tested GPT-4o-mini [28] to evaluate how effectively a large language model can predict actions from structured KG data. To interface with the model, we transformed each graph instance into a set of textual triples describing actions, objects, and their relations. We then used few-shot prompting [4] with structured reasoning, where the model was shown several examples that included intermediate steps to guide its inference process (see Figure2). This setup allowed the model to learn patterns from analogous cases and apply them to new graph segments at inference time. The code to replicate the results in this paper is available online.³

4. Results

Our experiments reveal a clear performance hierarchy across methods. For parent action prediction, frequency-based baselines (76–100% Hits@3) and GPT-4o-mini (78.73% Hits@1) outperformed graph-based models ($\leq 5.72\%$ Hits@5), suggesting that statistical patterns or high-level reasoning suffice for coarse task recognition. In contrast, sub-action prediction favored heuristic baselines (81.7% Hits@1), while graph models like RotatE (52.16% Hits@1) showed limited but non-trivial capability—likely due to their ability to model localized action transitions. Strikingly, LLMs collapsed to 13% accuracy here, exposing their weakness in fine-grained sequential reasoning.

²https://github.com/HRI-EU/BasicMemnet/tree/master/data/action_sequences

³https://github.com/marusta/Knowledge_Refinement

You are a knowledgeable assistant. Given a set of triples representing sub-actions, predict the most likely parent action.

Below are examples of parent actions and their corresponding sequences of sub-actions. Each triple represents a head, relation, and tail from the knowledge graph describing the parent action.

{examples}

What is the parent action for the following graph?

{test set triples}

Answer nothing but one of the 9 possible parent actions: cooking, cooking_with_bowls, pouring, wiping, cereals, hard_drive, free_hard_drive, hammering, sawing.

Figure 2. Prompt used for parent action prediction. It follows a few-shot format with structured reasoning. Parent action labels are taken directly from the dataset, where some tasks are object-based rather than action-based. The same naming is used in the prompt examples.

Category	Model	Parent Action Prediction			Subsequent Action Prediction		
		Hits@1	Hits@3	Hits@5	Hits@1	Hits@3	Hits@5
Baselines	Random	11%	33%	-	7.14%	-	-
	Baseline1	11%	76%	-	69.1%	-	-
	Baseline2	76%	100%	-	81.7%	-	-
KG Models	TransE	2.60%	4.90%	5.30%	0.00%	53.40%	64.86%
	TransR	3.25%	4.78%	5.72%	0.00%	31.51%	45.10%
	RotatE	3.88%	3.90%	3.92%	52.16%	57.67%	58.00%
	ComplEx	1.00%	2.00%	2.30%	8.56%	18.20%	24.18%
	DistMult	0.00%	0.00%	0.00%	33.68%	57.57%	58.84%
	ComplExLiteral	3.50%	3.80%	2.90%	14.49%	27.14%	33.91%
	DistMultLiteral	1.50%	3.00%	3.40%	5.40%	11.38%	16.10%
Other	GPT-4o-mini	78.73%	-	-	13.00%	-	-

Table 3. Parent action prediction and subsequent action prediction results

This failure aligns with the previous studies [14] noting that many embedding-based models, such as TransE and TransR, struggle with multistep dependencies, which is crucial for accurate parent action prediction. It also highlights another limitation: disconnected subgraphs in robotic tasks violate the dense-connectivity assumptions of conventional KG benchmarks.

5. Conclusion and Outlook

This work identifies robotic action prediction as a critical testbed for KGC methods, where standard LP techniques underperform due to disconnected task subgraphs and hierarchical dependencies. While baselines and LLMs offer pragmatic solutions for specific tasks, their limitations underscore the need for specialized graph representations that encode temporal, compositional, and weakly connected structures.

Future research should redefine evaluation protocols for LP in disconnected graphs, moving beyond random link masking, develop hybrid architectures that integrate baseline robustness with KG relational reasoning, and explore dynamic graph embeddings to capture action progression. By addressing these gaps, the field can unlock the KGC methods’ potential for real-world sequential decision-making.

References

- [1] Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Sahand Sharifzadeh, Volker Tresp, and Jens Lehmann. PyKEEN 1.0: A Python Library for Training and Evaluating Knowledge Graph Embeddings. *Journal of Machine Learning Research*, 22(82):1–6, 2021.
- [2] Michael Beetz, Rolf Dillmann, Martin Hanheide, Michael Kirchberg, Dimos Kulic, Thomas Löwe, Yutaka Maeda, Alexander Muelling, Bernhard Nebel, Wolfram Nebel, et al. Towards automated models of activities of daily life. *IEEE Robotics & Automation Magazine*, 17(1):27–40, 2010.
- [3] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc.
- [5] Borui Cai, Yong Xiang, Longxiang Gao, He Zhang, Yunfeng Li, and Jianxin Li. Temporal knowledge graph completion: A survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-2023*, page 6545–6553. International Joint Conferences on Artificial Intelligence Organization, August 2023.
- [6] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision. *International Journal of Computer Vision*, 130(1):33–55, 2022.
- [7] P. Das, C. Xu, R. F. Doell, and Corso J. J. A thousand frames in just a few words: Lingular description of videos through latent topics and sparse object stitching. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [8] Christian R. G. Dreher, Mirko Wächter, and Tamim Asfour. Learning object-action relations from bimanual human demonstration using graph networks. *IEEE Robotics and Automation Letters (RA-L)*, 5(1):187–194, 2020.
- [9] Julian Eggert, Jörg Deigmöller, Pavel Smirnov, Johane Takeuchi, and Andreas Richter. Memory net: Generalizable common-sense reasoning over real-world actions and objects. In *International Conference on Knowledge Engineering and Ontology Development*. Scitepress, November 2023.
- [10] Lisa Ehrlinger and Wolfram Wöb. Towards a definition of knowledge graphs. *SEMANTICS (Posters, Demos, SuCCESS)*, 48(1-4):2, 2016.
- [11] Sourav Garg, Niko Sünderhauf, Feras Dayoub, Douglas Morrison, Akansel Cosgun, Gustavo Carneiro, Qi Wu, Tat-Jun Chin, Ian Reid, Stephen Gould, et al. Semantics for robotic mapping, perception and interaction: A survey. *Foundations and Trends® in Robotics*, 8(1–2):1–224, 2020.
- [12] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D’amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge graphs. *ACM Comput. Surv.*, 54(4), July 2021.
- [13] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514, 2022.
- [14] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2):494–514, 2022.
- [15] Seyed Mehran Kazemi, Rishab Goel, Kshitij Jain, Ivan Kobyzev, Akshay Sethi, Peter Forsyth, and Pascal Poupart. Representation learning for dynamic graphs: A survey, 2020.
- [16] Yu Kong and Yun Fu. Human action recognition and prediction: A survey. *International Journal of Computer Vision*, 130(5):1366–1401, 2022.
- [17] Franz Krause, Tobias Weller, and Heiko Paulheim. On a generalized framework for time-aware knowledge graphs. In *Towards a Knowledge-Aware AI*, pages 69–74. IOS Press, 2022.

- [18] Agustinus Kristiadi, Mohammad Asif Khan, Denis Lukovnikov, Jens Lehmann, and Asja Fischer. Incorporating literals into knowledge graph embeddings. In *The Semantic Web–ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part I 18*, pages 347–363. Springer, 2019.
- [19] H. Kuehne, A. B. Arslan, and T. Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *Proceedings of Computer Vision and Pattern Recognition Conference (CVPR)*, 2014.
- [20] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-to-end generative framework for video segmentation and recognition. In *Proc. IEEE Winter Applications of Computer Vision Conference (WACV 16)*, Lake Placid, Mar 2016.
- [21] Ang Li, Meghana Thotakuri, David A. Ross, João Carreira, Alexander Vostroikov, and Andrew Zisserman. The ava-kinetics localized human actions video dataset, 2020.
- [22] Yin Li, Miao Liu, and James M. Rehg. In the eye of the beholder: Gaze and actions in first person video, 2020.
- [23] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. Learning entity and relation embeddings for knowledge graph completion. *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1), Feb. 2015.
- [24] Miao Liu, Siyu Tang, Yin Li, and James M Rehg. Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 704–721. Springer, 2020.
- [25] David N Nicholson and Casey S Greene. Constructing knowledge graphs and their biomedical applications. *Computational and structural biotechnology journal*, 18:1414–1428, 2020.
- [26] Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2016.
- [27] Alberto Olivares-Alarcos, Daniel Beßler, Alaa Khamis, Paulo Goncalves, Maki K. Habib, Julita Bermejo-Alonso, Marcos Barreto, Mohammed Diab, Jan Rosell, João Quintas, and et al. A review and comparison of ontology-based approaches to robot autonomy. *The Knowledge Engineering Review*, 34:e29, 2019.
- [28] OpenAI. Gpt-4o-mini, 2023.
- [29] Marcus Patel and Ethan Reynolds. Robotic perception and manipulation in unstructured environments. *International Journal on Mechanical Engineering and Robotics*, 13(1):16–22, 2024.
- [30] Heiko Paulheim. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic web*, 8(3):489–508, 2016.
- [31] David Paulius, Yongqiang Huang, Roger Milton, William D. Buchanan, Jeanine Sam, and Yu Sun. Functional object-oriented network for manipulation learning. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, page 2655–2662. IEEE Press, 2016.
- [32] Francesca Pistilli and Giuseppe Averta. Graph learning in robotics: a survey. *IEEE Access*, 11:112664–112681, 2023.
- [33] Andrea Rossi, Denilson Barbosa, Donatella Firmani, Antonio Matinata, and Paolo Merialdo. Knowledge graph embedding for link prediction: A comparative analysis. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(2):1–49, 2021.
- [34] Yaoxian Song, Penglei Sun, Haoyu Liu, Zhixu Li, Wei Song, Yanghua Xiao, and Xiaofang Zhou. Scene-driven multimodal knowledge graph construction for embodied ai. *IEEE Trans. on Knowl. and Data Eng.*, 36(11):6962–6976, November 2024.
- [35] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space, 2019.
- [36] Moritz Tenorth and Michael Beetz. Knowrob: A knowledge processing infrastructure for cognition-enabled robots. *The International Journal of Robotics Research*, 32(5):566–590, 2013.
- [37] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction, 2016.
- [38] Cordula Vesper. How to support action prediction: Evidence from human coordination tasks. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pages 655–659. IEEE, 2014.
- [39] Markus Waibel, Michael Beetz, Javier Civera, Raffaello D’Andrea, Jos Elfring, Dorian Gálvez-López, Kai Häussermann, Rob Janssen, J.M.M. Montiel, Alexander Perzylo, Björn Schießle, Moritz Tenorth, Oliver Zweigle, and René Van De Molengraft. Roboearth. *IEEE Robotics & Automation Magazine*,

- 18(2):69–82, 2011.
- [40] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, page 950–958, New York, NY, USA, 2019. Association for Computing Machinery.
 - [41] Bishan Yang, Wen tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases, 2015.
 - [42] Andrea Maria Zanchettin, Andrea Casalino, Luigi Piroddi, and Paolo Rocco. Prediction of human activity patterns for human–robot collaborative assembly tasks. *IEEE Transactions on Industrial Informatics*, 15(7):3934–3942, 2018.
 - [43] Shengchen Zhang, Zixuan Wang, Chaoran Chen, Yi Dai, Lyumanshan Ye, and Xiaohua Sun. Patterns for representing knowledge graphs to communicate situational knowledge of service robots. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery.