

HTEKG: A Human-Trait-Enhanced Literary Knowledge Graph with Language Model Evaluation

Sreejith Sudhir Kalathil¹, Tian Li¹, Hang Dong² and Huizhi Liang¹

¹Newcastle University, Newcastle upon Tyne, U.K.

²University of Exeter, Exeter, U.K.

Keywords: Knowledge Graph, Literary Analysis, Language Model.

Abstract: Knowledge Graphs (KGs) are a crucial component of Artificial Intelligence (AI) systems, enhancing AI's capabilities in literary analysis. However, traditional KG designs in this field have focused more on events, often ignoring character information. To tackle this issue, we created a comprehensive Human-Trait-Enhanced Knowledge Graph, HTEKG, which combines past event-centered KGs with general human traits. The HTEKG enhances query capabilities by mapping the complex relationships and traits of literary characters, thereby providing more accurate and context-relevant information. We tested our HTEKG on three typical literary comprehension methods: traditional Cypher query, integration with a BERT classifier, and integration with GPT-4, demonstrating its effectiveness in literary analysis and its adaptability to different language models.

1 INTRODUCTION

The advent of Artificial Intelligence (AI) has significantly revolutionized our ability to analyze, interpret, and interact with human literature, offering several powerful tools for studying vast volumes of literary data, including summarization, topic and relationship mining, entity extraction, information structuring, and retrieval. Meanwhile, these technologies facilitate researchers in achieving advancements in domains such as human-computer interaction, role-play generation, and creative writing assistance. However, despite its prosperity, AI systems still encounter challenges in approaching the human level. This is because current NLP technology mainly focuses on information and knowledge while overlooking personal trait and relation aspects. For example, prior works (Van Hage and Ceolin, 2013; Lombardo et al., 2018; Kozaki et al., 2023; P. Wilton, 2013; Khan et al., 2016; Yeh, 2017; Lisena et al., 2023) have concentrated on author-centric and event-centric knowledge points, such as authorship, content summaries, genre information, story elements, event timelines, and their interrelations. The forced application of knowledge-centered schemas to character analysis inevitably leads to the loss of nuanced information about individuals and the complex relationships between them (Ugai et al., 2024; Ugai, 2023). Consequently, it constrains the AI potential in tasks such as character analysis and

role-playing, hindering AI's application in fields like creative writing, game development, AI in education, and person-centered recommendations.

To address these issues, this paper aims to construct a comprehensive, Human-Trait-Enhanced Knowledge Graph with a new backbone ontology schema to depict human and contextual attributes across various literary pieces. We name the knowledge graph built upon this ontology HTEKG. By extracting relevant entities and information from various literary datasets (such as Project Gutenberg¹ and Goodreads²) according to our proposed ontology, we build HTEKG on the Neo4j platform. To evaluate the effectiveness of our system for character understanding and its compatibility with cutting-edge NLP models, we employed three different methods, Neo4j Cypher, BERT classification (Devlin et al., 2018), and GPT-4 (Achiam et al., 2023), to assess HTEKG on multiple dimensions.

Our key contributions are as follows:

- We pioneered a Human-Trait-Enhanced Knowledge Graph schema and extraction process, incorporating character attributes, emotions, and relational dynamics into the knowledge graph in a standardized manner. This endows AI frameworks with more comprehensive literary character

¹<https://www.gutenberg.org>

²<https://www.goodreads.com>

analysis capabilities, and the standardization and generality of our approach ensure its scalability to the vast amount of literature.

- Our detailed evaluation on integrating HTEKG with BERT and GPT-4 showcases not only its correctness and relatedness (Nguyen et al., 2023) in character understanding but also its flexibility in integrating with the language models.

2 RELATED WORK

We define Knowledge Graphs (KGs) based on the various definitions in the field (Ehrlinger and Wöß, 2016), as schema (or a backbone *ontology*) of classes with their data instances. The schema provides high-level terms (corresponding to TBox in an OWL ontology (Grau et al., 2008)) and the facts provide instance level information (corresponding to ABox). A KG can be represented as a set of triples (s, p, o) of subject, predicate, and object.

KGs have been extensively applied in various narrative text scenarios to manage information and assist tasks such as reading comprehension, logical reasoning, and literary creation (Hitzler and Janowicz, 2013; Lehmann, 2015; Lee and Chang, 2019). There is a substantial amount of prior work on the design of ontologies for narrative text scenarios. The Simple Event Model (SEM) (Van Hage and Ceolin, 2013) was one of the earliest attempts to construct a unified ontology for narrative texts, focusing on the chronological order of events and forming their ontology around temporal sequences. However, their work neglected the characterization of people, thus omitting attributes like “who and how”. The work (Lombardo et al., 2018) proposed another ontology, Drammar, that included human attributes, but it was fiction-specific, lacking standardization and generalizability. The KG Reasoning Challenge (Kozaki et al., 2023) provided a KG for Sherlock Holmes novels, but its characterization of characters was too simplistic, missing information such as character motivations, opportunities, and means of crime. Additionally, many ontology works for narrative content (P. Wilton, 2013; Khan et al., 2016; Yeh, 2017; Lisena et al., 2023) opted for event-centric construction schemes. These studies focus on addressing the challenges of constructing nonlinear narratives using KGs rather than on character and relationship depiction.

In terms of KGs usage, early approaches mainly involved retrieval reasoning and embedding techniques. In the KG Reasoning Challenge, Ugai (Ugai et al., 2024; Ugai, 2023) constructed a KG

encompassing character motives, opportunities, and methods, coupling it with an event-centered KG to achieve interpretable crime predictions. Kurokawa (Kurokawa, 2021) used various KG embeddings to perform crime prediction and summarization through link prediction. Nguyen (Nguyen et al., 2023) employed BERT classification to test their biological KG. In recent years, with the rise of large language models (LLMs), graph-based retrieval-augmented generation (RAG) systems that integrate KGs and LLMs have become increasingly popular. Ashby (Ashby et al., 2023) utilizes KGs to assist LLMs in procedural content generation. Mindmap (Wen et al., 2023) and GraphRAG (Edge et al., 2024) constructs multi-level natural language summaries (also named “community”) for KG retrieval to facilitate the use of KG information by LLMs. The core objective of these approaches is to leverage the topological information of KGs to represent complex relationships, thereby enhancing the effectiveness of the final system. To demonstrate the adaptability of our ontology to different approaches, we selected one representative approach from each of the three categories above (Neo4j Cypher, BERT-based classification, and GPT-4 API) and coupled it with our HTEKG for evaluation.

3 METHODOLOGY

In this section, we elaborate on the following three topics to create our HTEKG including schema design, construction, and evaluation:

- **Schema Design of the HTEKG.** How to design a Human-Trait-Enhanced backbone ontology to represent character information, particularly the complex relationships and interaction features among characters.
- **Construction of the HTEKG from Texts.** How to utilize NLP techniques to extract the corresponding entities and populate the final HTEKG with the backbone ontology. This involves named entity recognition (NER) and subject-verb-object (SVO) methods for character identification, geographic detection, and relationships.
- **Evaluation of the HTEKG.** How to couple our HTEKG with different language models (BERT and GPT-4) to give predictions of detailed character attributes and relationships, or generations for evaluation purposes.

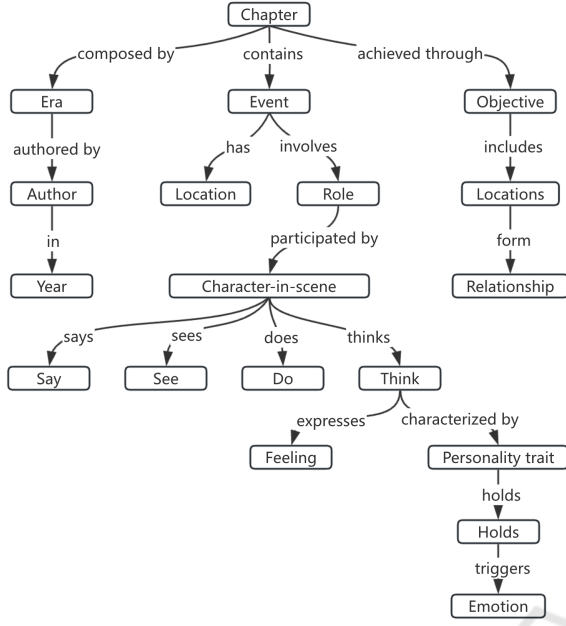


Figure 1: Human-Trait-Enhanced KG schema.

3.1 Schema Design of the HTEKG

To build a unified literary ontology, we integrated the Simple Event Model (SEM) (Van Hage and Celin, 2013; Gottschalk and Demidova, 2018), Empathy Knowledge Representation (Pileggi, 2021), ABC Model of Personality (Ellis, 1991), and theory of mind for cognitive KG (Wu et al., 2023), and make adjustments to better depict character traits and relationships. Our goal is to provide a hierarchical and interconnected framework containing semantically rich character attributes and narrative elements, thereby promoting a deeper and more comprehensive literature analysis through the KG.

The overall schema is illustrated in Figure 1. It has several key design objectives. First, it provides enhanced character and narrative integration, which includes the following features:

- **Chapter and Contextual Data.** The schema begins with the class “Chapter”, which then contextualizes the narrative by including “Era”, “Author”, and “Year”. This foundational layer sets the stage for a deeper understanding of the narrative and its historical or stylistic settings.
- **Events and Settings.** Each “Event” within a chapter includes the “Location” where it unfolds, providing spatial context that influences character interactions and developments.
- **Role Dynamics.** The roles that characters play within these events are depicted as the flow, changing from one event to another, reflecting the

evolving nature of their interactions and the significance within the narrative.

Second, the backbone ontology is designed to provide detailed character-in-scene depiction, as listed below:

- **Actions and Beliefs.** The “Character-in-scene” component captures what characters see, say, and do, along with their thoughts and beliefs. These elements are crucial for illustrating “Personality Traits” and how these traits manifest in different situations (Lehmann, 2015).
- **Emotions.** Characters’ emotions are intricately linked to events, showcasing how different scenarios influence their emotional responses. This dynamic portrayal helps in understanding how characters react under varying circumstances.

Third, this schema displays **linkages of several narrative elements**, including the linking between the chapter and its main objectives and the relationships among characters within that chapter. Such linkages also provide more in-depth social fabric visualization based on the interactions between characters, such as friendships, rivalries, alliances, and familial bonds. This feature adds ontological significance to the characters, enriching the potential to depict their human traits.

These three features are vital to illustrate how characters’ goals and interactions drive the narrative forward and affect their development, and is something that previous work has overlooked.

3.2 Construction of the HTEKG

This section introduces two parts: (i) the extended settings used to construct the HTEKG; (ii) the NLP techniques employed for various information extraction.

3.2.1 Extended Settings to Construct HTEKG

For the definition of nodes, each of them has multiple attributes, such as character and emotional states. For example:

- **Character Nodes.** Attributes include roles (“HAS_ROLE”) and traits (“HAS_TRAIT”), providing a dynamic representation of character development.
- **Scene Nodes.** Attributes related to characters and settings, reflecting the spatial dynamics and emotional textures of the narrative.

For the definition of edges, we use relationships such as “INCLUDES_SCENE” (between “Chapter” and “Scene” nodes) or “FEATURES_IN” (between “Chapter” nodes) to connect different nodes.

3.2.2 Techniques for Information Extraction

Before performing all extraction steps, we preprocessed the data by removing disturbing symbols and irrelevant information (such as project introductions and copyright statements) that could interfere with subsequent operations. By identifying specific chapter markers, we divided the data into different chapters. This approach facilitates targeted extraction of the content, themes, and objectives within each chapter in the subsequent steps (Siahaan et al., 2023).

For entities such as names, locations, and dates, we utilized SpaCy with their official English NER model for extraction. For character traits, emotions, and personality characteristics, we first located them using SpaCy's tags and then matched and identified these features using a predefined sentiment lexicon.

To extract character relationships, we employ another lexicon dictionary (a small number of pre-prepared words and use WordNet (Miller, 1995) to enumerate all their synonyms) to match and map the results to predefined relationships. This approach mitigates the omission caused by unconventional expressions of some relationships.

3.3 Evaluation of the HTEKG

The evaluation compares the results of the querying of the HTEKG (e.g. with a KG querying language: Cypher) with pre-prepared answers. We calculated the accuracy (Acc), precision (Prec), recall (Rec), and F1 score of the retrieval results, thus reflecting the effectiveness of our HTEKG in literary analysis.

We also tested fusing HTEKG with two language models: GPT-4 and BERT (Nguyen et al., 2023). This provides further insights into our HTEKG's adaptation to cutting-edge models. Their processes are introduced in the following sections.

3.3.1 Integration with GPT-4

For the integration with GPT-4, the overall process is divided into three steps:

1. We concatenated the natural language requirement query and the structural description of KG and input them into GPT4 to generate the corresponding Cypher query. An example is shown in Table 1.
2. We execute the Cypher query on Neo4j to retrieve the corresponding search results.
3. We concatenate the query and the Cypher query results and input them into GPT-4 to obtain the final output. An example is shown in Table 2.

We assess the quality of the generated content through manual annotation. For information of the test samples and the GPT version, please refer to Section 4.

3.3.2 Integration with BERT

For the integration with BERT, the overall process involves training and evaluation phases, as we convert it into multiple classification tasks and achieve it by fine-tuning BERT. The using of fine-tuned BERT for KG evaluation is inspired by the previous work in (Nguyen et al., 2023). Four classifiers are prepared, including Traits Classifier, Relationships Classifier, Emotions Classifier, and Events Classifier. The task of each classifier is a binary classification task. The meaning of the label is defined as follows:

- **Correct Information (Label = 1).** This label signifies that the extracted attribute (such as traits, relationships, emotions, or events) accurately matches the information from the novel. For example, if the classifier correctly identifies "Sherlock Holmes" as "intelligent", it is labeled as correct.
- **Incorrect Information (Label = 0).** This label indicates that the attribute is either incorrect or inconsistent with the known data. For instance, if the dataset incorrectly states that "Sherlock Holmes" is "fearful", it is labeled as incorrect.

After manually collecting the samples required for each task, the entire dataset is divided into two parts: a training set, used for fine-tuning the classifiers, and a test set, used to evaluate their performance. For meta information and statistical information about the dataset, please refer to Section 4.

4 EXPERIMENTS

We used Neo4j as the platform to construct our HTEKG. It leverages a graph model to represent and navigate relationships between data points, making it ideal for applications that require connected data insights. Meanwhile, it employs Cypher, a powerful and expressive query language specifically designed for working with graph data. Additionally, it has several build-in tools to visualize the database.

We used the novel "The Hound of the Baskervilles" from the Gutenberg dataset to construct the HTEKG. This knowledge graph contains 239 nodes and 683 relationships. The statistics of entities extracted during the NER phase are shown in

Table 1: The sample to prompt GPT-4 to convert natural language query plus graph structure information into Cypher query.

Query	Graph Structure Description	Response
Given the context of ‘The Hound of the Baskervilles’ by Sir Arthur Conan Doyle, generate a Cypher query to find the characters who are friends of Dr. Watson. Return the query in JSON format.	<p>Nodes: Character, Book, Chapter, Location, Event, Scene</p> <p>Character nodes have properties: name, ...</p> <p>Book nodes have properties: title, ...</p> <p>Relationships:</p> <p>Character nodes are connected to Book, ...</p> <p>Character nodes can have relationships with each other (e.g., FRIENDS_WITH, ...).</p> <p>...</p>	<pre>MATCH (c:Character)- [r:FRIEND]->(w:Character {name: 'Dr. Watson'}) RETURN c.name AS friend_name</pre>

Table 2: The sample to prompt GPT-4 to generate the answer based on query and KG information.

Query	Cypher Result	Response
Describe the relationship between Dr. Watson and Sherlock Holmes.	<pre>{'characterProfile': {'name': 'Sherlock Holmes', 'roles': [('role': 'Detective', 'description': 'Holmes uses his remarkable deductive reasoning skills to solve complex and mysterious cases.'], ('role': 'Friend', 'description': 'Holmes is a close friend and flatmate to Dr. John Watson, often working together on various cases.'], 'traits': ['Intellectual', 'Observant', 'Analytical', 'Eccentric', 'Unemotional'], 'emotionalExperiences': ['experience: 'Frustration', 'description': 'Holmes often shows frustration when people fail to see what he considers obvious clues.'], ...</pre>	Dr. Watson and Sherlock Holmes share a deep, rich friendship and an incredibly strong professional relationship as well. Introduced in Arthur Conan Doyle’s novel, “A Study in Scarlet,” Dr. John Watson meets Sherlock Holmes when both are looking for affordable lodgings, leading to them sharing rooms at 221B Baker Street...




Figure 2: Statistics of entities for “The Hound of the Baskervilles”.

Figure 2. Figures 3 and 4 illustrate how the HTEKG represents Jack Stapleton and the relationship between Dr. John Watson and Dr. James Mortimer.

For the traditional Cypher queries and BERT classification scenarios, we prepared 500 samples across 4 categories, with no fewer than 100 samples per category. Since the BERT classifiers require training, the dataset was split into a 7:3 ratio for training and evaluation. The evaluation for both Cypher queries and BERT classifiers was conducted on the evaluation set. The detailed statistics are shown in Table 3.

For training the BERT model, we used Bert-base-uncased as our base model and fine-tuned it to create the four classifiers as needed.

For the dataset used to test the integration with

Table 3: The dataset statistics for Cypher queries and BERT classifiers.

Classifier	Train	Evaluation	Total
Character Traits	105	45	150
Relationships	70	30	100
Emotions	84	36	120
Events	91	39	130

GPT-4, we manually prepared 24 queries for evaluation. For the GPT-4 version, we used the GPT-4-0613 model. This is the standard version provided by OpenAI, supporting a token length of 8K.

5 RESULTS

We report results based on the three evaluation strategies, comparing KG query (with Cypher language) to predefined answers, comparing fine-tuned BERT results with the HTEKG, and integrating HTEKG with prompting GPT-4.

5.1 Results of Querying HTEKG

Table 4 demonstrates the results of using Cypher queries to retrieve relevant characters, emotions, relationships, and events. It showcases the HTEKG’s ability to structurally extract and present complex character dynamics from literary texts.

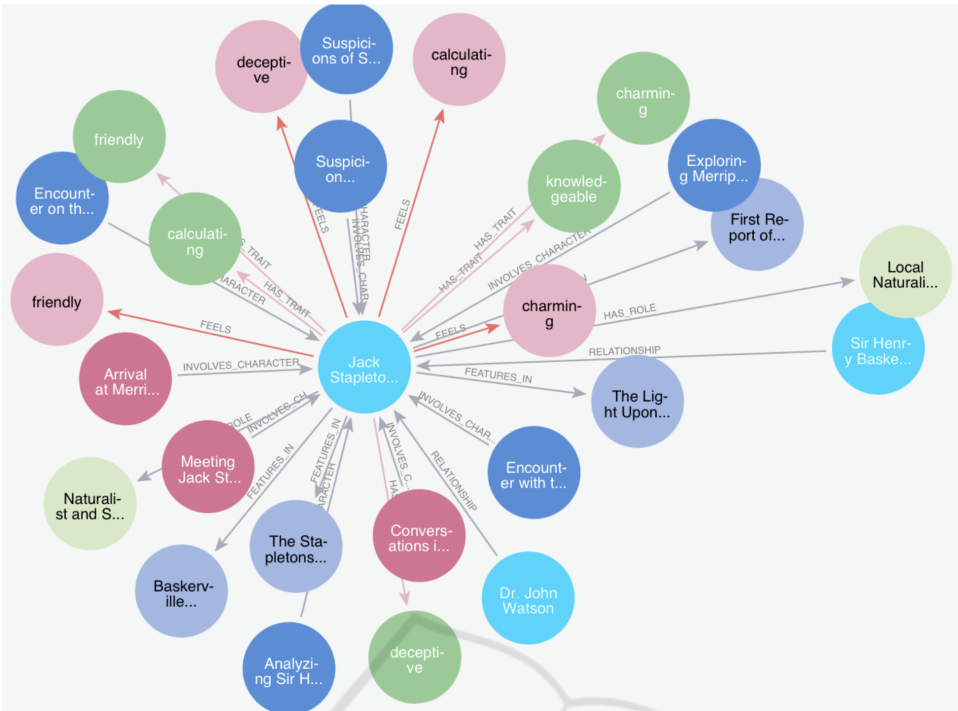


Figure 3: The demonstration subgraph of Jack Stapleton extracted from the HTEKG.

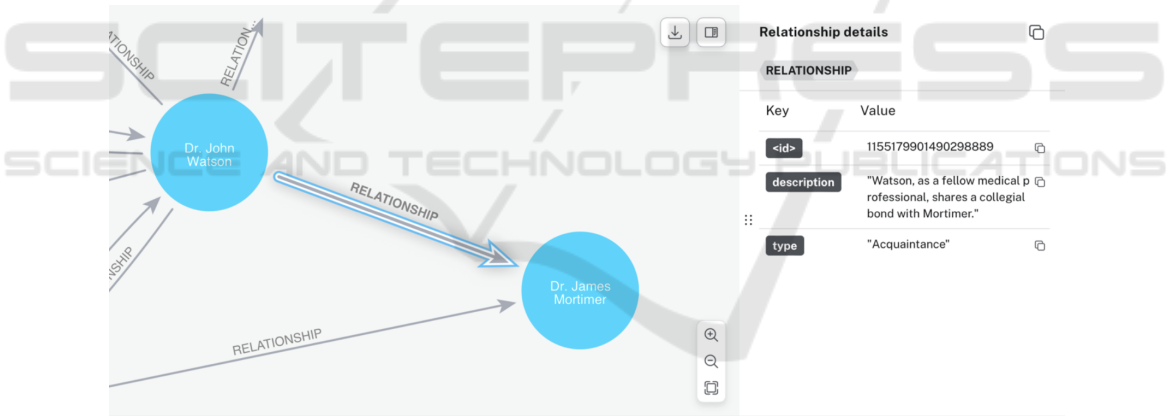


Figure 4: The demonstration of the “acquaintance” relationship between Dr. John Watson and Dr. James Mortimer.

Table 4: Results of querying HTEKG using Cypher.

Category	Acc	Prec	Rec	F1
Character Traits	0.91	0.90	0.91	0.90
Relationships	0.77	0.74	0.76	0.75
Emotions	0.86	0.84	0.83	0.83
Events	0.87	0.85	0.86	0.85

5.2 Result of Integrating HTEKG with BERT

Table 6 shows the performance of integrating HTEKG with BERT in predicting different attributes. It excels in predicting traits. However, it performs moderately

in predicting attributes with complex relations, such as character relationships, emotions, and the associated events.

5.3 Result of Integrating HTEKG with GPT-4

Table 5 displays a full GPT-4 respond to the query based on the HTEKG content, and Table 7 provides metrics proving that the HTEKG integrates well with GPT-4. It can be observed that even when concatenating long Cypher results, GPT-4 can still understand the content and extract relevant information to gener-

Table 5: An example of GPT-4’s response.

Input	Response
List the implications of this narrative element in bullet points: Sir Charles Baskerville’s death. + \${Cypher Result}	<p>The death of Sir Charles Baskerville kick-starts the main storyline of the novel “The Hound of the Baskervilles” by Arthur Conan Doyle involving a legendary beast supposedly haunting the Baskerville family.</p> <ul style="list-style-type: none"> - It brings into focus the supposed curse on the Baskerville family which becomes a subject of investigation and interest. - Sir Charles’s death leads to the introduction of Sherlock Holmes and Dr. John Watson, the main characters, who are brought in to investigate the mysterious circumstances surrounding his death. - After his death, the heir apparent, Sir Henry Baskerville, is introduced, which adds another layer of suspense and intrigue to the plot. - It leads to the exploration of the Baskerville Hall and its surrounding moor.

Table 6: Results of BERT classifiers.

Category	Acc	Prec	Rec	F1
Character Traits	0.87	0.85	0.84	0.84
Relationships	0.67	0.65	0.66	0.65
Emotions	0.75	0.73	0.74	0.73
Events	0.77	0.76	0.77	0.77

Table 7: Results using GPT-4-0613.

Query Method	Accuracy
GPT-4	0.63
+ HTEKG	0.85

ate context-rich and correct responses. The generated content includes a great deal of detail, which is not visible when querying GPT-4 without the HTEKG. The results highlight the HTEKG’s potential to improve AI tools for literary analysis.

5.4 Discussion

Beyond the evaluation metrics presented above, we also observed that integrating HTEKG with GPT-4 enhanced GPT-4’s ability to provide more detailed and nuanced explanations. This opens up new avenues for interpretability in future research, allowing researchers to monitor the internal thought processes of GPT-4 more effectively. Also, GPT-4 excels at using HTEKG to generate coherent predictions when dealing with complex narrative elements such as relationships. In contrast, BERT may give contradictory results when predicting the relationship between two people using different person’s subgraphs.

6 CONCLUSION

In this work, we propose a novel standardized Human-Trait-Enhanced Knowledge Graph, HTEKG. Our ontology highlights the nuanced interplay of character traits and relationships, showcasing the potential of KGs to capture complex human interactions within the literature. We evaluated HTEKG in three scenarios: Cypher query, integration with BERT, and integration with GPT-4. Our evaluation demonstrates the role of HTEKG in literary understanding and its flexibility in coupling with different AI frameworks.

However, this work has the following limitations:

- *Comprehensive Knowledge Coverage.* This work primarily focuses on integrating general human attributes with past event-centered ontology, lacking coverage in areas such as temporal knowledge. Additionally, there is a lack of specific research on idioms and metaphors in literary works, which may result in some knowledge missing.
- *Evaluation Method Limitations.* The three evaluation scenarios used in this work are basic and general practices compared to the real-world application. More evaluation methods need to be attempted to verify the effectiveness of HTEKG with various cutting-edge technologies. Moreover, the current HTEKG evaluation lacks standardized metrics, so we devised different evaluation schemes based on usage scenarios. This must be tackled in the future.

For future work, we propose the following directions:

- *More Advanced Knowledge Extraction Techniques.* We aim to introduce more advanced and automated technologies, leveraging graph under-

standing AI frameworks, to standardize and accelerate the knowledge extraction process.

- *Broader Corpus Coverage.* We plan to build our HTEKG on a larger scale corpus.
- *Clearer Downstream Application Scenarios.* Besides analysis, we will also explore using the HTEKG to assist in tasks such as role-playing and procedural content generation. For instance, in role-playing scenarios, current systems that combine LLMs with agent-based design are highly disorganized, with each implementation using its own unique memory design. However, since our ontology encompasses detailed role information (including states, goals, emotions, and past actions), our HTEKG can be applied to standardize the design of agent systems in this domain, thereby enhancing and scaling up the role-playing capabilities of existing LLMs.

REFERENCES

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ashby, T., Webb, B. K., Knapp, G., Searle, J., and Fulda, N. (2023). Personalized quest and dialogue generation in role-playing games: A knowledge graph-and language model-based approach. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–20.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., and Larson, J. (2024). From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Ehrlinger, L. and Wöß, W. (2016). Towards a definition of knowledge graphs. *SEMANTiCS (Posters, Demos, SuCCESS)*, 48(1-4):2.
- Ellis, A. (1991). The revised abc's of rational-emotive therapy (ret). *Journal of rational-emotive and cognitive-behavior therapy*, 9(3):139–172.
- Gottschalk, S. and Demidova, E. (2018). Eventkg: A multilingual event-centric temporal knowledge graph. In *The Semantic Web: 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, Proceedings 15*, pages 272–287. Springer.
- Grau, B. C., Horrocks, I., Motik, B., Parsia, B., Patel-Schneider, P., and Sattler, U. (2008). Owl 2: The next step for owl. *Journal of Web Semantics*, 6(4):309–322. Semantic Web Challenge 2006/2007.
- Hitzler, P. and Janowicz, K. (2013). Linked data, artificial intelligence, and the semantic web. *AI Magazine*, 34(2):63–72.
- Khan, A. F., Bellandi, A., Benotto, G., Frontini, F., Giovannetti, E., and Reboul, M. (2016). Leveraging a narrative ontology to query a literary text. In *7th Workshop on Computational Models of Narrative (CMN 2016)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.
- Kozaki, K., Egami, S., Matsushita, K., Ugai, T., Kawamura, T., and Fukuda, K. (2023). Datasets of mystery stories for knowledge graph reasoning challenge. In *ESWC Workshops*.
- Kurokawa, M. (2021). Explainable knowledge reasoning framework using multiple knowledge graph embedding. In *Proceedings of the 10th International Joint Conference on Knowledge Graphs*, pages 172–176.
- Lee, Y. and Chang, H. (2019). The role of ontology in understanding human relations. *Cognitive Science Quarterly*, 39(1):45–65.
- Lehmann, J. (2015). Dbpedia – a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195.
- Lisena, P., Schwabe, D., van Erp, M., Leemans, I., Tullett, W., Troncy, R., Ehrich, S. C., and Ehrhart, T. (2023). An ontology for creating hypermedia stories over knowledge graphs. In *SWODCH 2023, Workshop on Semantic Web and Ontology Design for Cultural Heritage*.
- Lombardo, V., Damiano, R., and Pizzo, A. (2018). Drammar: A comprehensive ontological resource on drama. In *The Semantic Web—ISWC 2018: 17th International Semantic Web Conference, Monterey, CA, USA, October 8–12, 2018, Proceedings, Part II 17*, pages 103–118. Springer.
- Miller, G. (1995). Wordnet: A lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Nguyen, H., Chen, H., Chen, J., Kargozari, K., and Ding, J. (2023). Construction and evaluation of a domain-specific knowledge graph for knowledge discovery.
- P. Wilton, J. Tarling, J. M. (2013). Storyline ontology - an ontology to represent news storylines.
- Pileggi, S. F. (2021). Knowledge interoperability and reuse in empathy mapping: An ontological approach. *Expert Systems with Applications*, 180:115065.
- Siahaan, D., Raharjana, I., and Fatichah, C. (2023). User story extraction from natural language for requirements elicitation: Identify software-related information from online news. *Information and Software Technology*, 158:107195.
- Ugai, T. (2023). Criminal investigation with augmented ontology and link prediction. In *2023 IEEE 17th International Conference on Semantic Computing (ICSC)*, pages 288–289. IEEE.
- Ugai, T., Koyanagi, Y., and Nishino, F. (2024). A logical approach to criminal case investigation. *arXiv preprint arXiv:2402.08284*.
- Van Hage, W. R. and Ceolin, D. (2013). The simple event model. In *Situation awareness with systems of systems*, pages 149–169. Springer.

- Wen, Y., Wang, Z., and Sun, J. (2023). Mindmap: Knowledge graph prompting sparks graph of thoughts in large language models. *arXiv preprint arXiv:2308.09729*.
- Wu, J., Chen, Z., Deng, J., Sabour, S., and Huang, M. (2023). Coke: A cognitive knowledge graph for machine theory of mind. *arXiv preprint arXiv:2305.05390*.
- Yeh, J.-h. (2017). Storyteller: An event-based story ontology composition system for biographical history. In *2017 International Conference on Applied System Innovation (ICASI)*, pages 1934–1937. IEEE.

