

SymbolicThought: Integrating Language Models and Symbolic Reasoning for Consistent and Interpretable Human Relationship Understanding

Runcong Zhao^{1*}, Qinglin Zhu^{1*}, Hainiu Xu¹, Bin Liang², Lin Gui¹, Yulan He^{1,3}

¹King’s College London, ²MoE Lab, CUHK, ³The Alan Turing Institute
{runcong.zhao, qinglin.1.zhu, hainiu.xu}@kcl.ac.uk, bin.liang@cuhk.edu.hk
{lin.1.gui, yulan.he}@kcl.ac.uk

Abstract

Understanding character relationships is essential for interpreting complex narratives and conducting socially grounded AI research. However, manual annotation is time-consuming and low in coverage, while large language models (LLMs) often produce hallucinated or logically inconsistent outputs. We present *SymbolicThought*, a human-in-the-loop framework that combines LLM-based extraction with symbolic reasoning. The system constructs editable character relationship graphs, refines them using seven types of logical constraints, and enables real-time validation and conflict resolution through an interactive interface. To support logical supervision and explainable social analysis, we release a dataset of 160 interpersonal relationships with corresponding logical structures. Experiments show that *SymbolicThought* improves annotation accuracy and consistency while significantly reducing time cost, offering a practical tool for narrative understanding, explainable AI, and LLM evaluation. The source code and dataset are publicly available on GitHub¹. Updated resources (e.g., screen-cast and demo links) will be maintained in the GitHub repository.

1 Introduction

Recent advancements in artificial intelligence, particularly the emergence of large language models (LLMs), have reignited discussions around the potential for artificial general intelligence. The increasingly human-like capabilities of AI are also drawing significant attention from social science researchers (Xu et al., 2024). One of the most critical functions in social science research is understanding the relationships between individuals. This understanding is foundational to various disciplines, including journalism, legal document analysis, historical studies, and literary interpretation.

Many existing studies suggest that social relationships can be modeled as compact graphs, also known as small-world networks. These graphs are characterised by a high clustering coefficient and short path lengths. Consequently, such networks typically contain **small but densely connected subgraphs** (communities), where nodes within a community are highly likely to be connected. Connections between different communities are often **facilitated by a few key nodes** (individuals).

However, constructing such graphs solely from plain text descriptions presents a significant challenge. This is because the dense local structure of social communities requires a method capable of capturing most intra-community relationships, while the accurate identification of key inter-community connections demands high precision in detecting critical relationships. Traditional methods rely on either **crowdsourced annotations** or automated approaches based on LLMs. Both methods exhibit substantial limitations.

On one hand, **human annotation** is generally accurate but demands significant cognitive effort and time investment, which usually lead to **low coverage** in the annotation. In particular, essential cross-sentence and long-distance connections are frequently missed or inaccurately recorded (Zhao et al., 2024; Sancheti and Rudinger, 2025). On the other hand, **LLM-based methods** provide scalability and efficiency but suffer from issues such as hallucinations, semantic inconsistency, and fragmented narrative representations (Liu et al., 2023; Guo et al., 2023; West et al., 2023; Yuan et al., 2024), which often fail to predict key relations precisely (Wang et al., 2024).

To address these challenges, we propose *SymbolicThought*, a hybrid framework designed for extracting character relationship graphs from narrative texts, combining the scalability and reasoning ability of LLMs with the precision and interpretability provided by human annotators. *Sym-*

*Equal contribution.

¹<https://github.com/BLPXSPG/SymbolicThought>

SymbolicThought integrates an LLM-based extraction mechanism with **interactive human-in-the-loop verification** supported by a novel **symbolic reasoning module** with an editable social knowledge base. Crucially, *SymbolicThought* leverages an intuitive visual interface that highlights relevant textual evidence and systematically detects logical inconsistencies or conflicts, significantly enhancing annotators’ focus and efficiency. By automatically surfacing critical information and reducing annotators’ cognitive load in searching through extensive narratives, our system empowers humans to direct their attention toward complex judgments and nuanced decisions, which are tasks that current fully automated systems are still unable to handle. Through this structured, human-centered annotation pipeline, *SymbolicThought* provides a solution that is:

- **Interpretable**, where the symbolic reasoning driven by the LLM is easily traceable, localisable, and understandable by users;
- **Extendable**, where users can edit the rules underlying the symbolic reasoning;
- **Complete**, where the automatic annotation framework is designed to generate all possible relation graphs for a given character list;

This enables the construction of an accurate and comprehensive static character relationship graphs, laying a strong foundation for advanced narrative analysis and socially aware AI tools across the humanities and social sciences.

2 Architecture of *SymbolicThought*

Given a narrative document \mathcal{D} , the goal of *SymbolicThought* is to construct a character relationship graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of unique, verified character nodes and \mathcal{E} is the set of semantically and logically consistent relationships between them. The pipeline consists of two main stages: (1) **Character Extraction**: An LLM-based module $\mathcal{F}_{\text{char}}$ proposes candidate entities from \mathcal{D} , which are then verified, grouped, or corrected by annotators to yield the finalised node set \mathcal{V} . (2) **Relationship Extraction**: A downstream extraction module \mathcal{F}_{ext} predicts pairwise relationships conditioned on both \mathcal{D} and \mathcal{V} . These predictions are refined by a symbolic module \mathcal{F}_{ref} and further reviewed by annotators to produce the final edge set \mathcal{E} .

2.1 Character Extraction

This step corresponds to the lower-left region of Figure 1. The goal is to identify a complete and clean set of character entities \mathcal{V} from the input narrative document \mathcal{D} , which serves as the foundation for downstream relation extraction.

Explorer: Node Candidate Generation. We run the character extractor $\mathcal{F}_{\text{char}}$ on the corpus \mathcal{D} independently n_c times, each with temperature sampling to broaden coverage. Mentions that recur in at least τ_c runs are retained, yielding a consensus set of candidate nodes:

$$\mathcal{V}_{\text{raw}} = \left\{ c \mid \sum_{i=1}^{n_c} \mathbb{I}[c \in \mathcal{F}_{\text{char}}^{(i)}(\mathcal{D})] \geq \tau_c \right\}.$$

This frequency filter maximises recall while suppressing spurious extractions, providing a high-coverage node list for subsequent verification.

Verifier: Human Verification. \mathcal{V}_{raw} is displayed in a front-end interface. Annotators can accept or reject suggestions, highlight new spans, or add missing entries. Mentions shift from **red** (unconfirmed) to **green** (confirmed), providing instant feedback. To improve coherence, our interface supports two editing operations: (1) Alias Resolution: Many characters appear under multiple aliases (e.g., “Elizabeth”, “Her Majesty”, or “Queen Elizabeth II”). Annotators can merge these into a single canonical entity via drag-and-drop operations. (2) Entity Disambiguation: In contrast, distinct characters may share the same surface name (e.g., “Alexandre Dumas” may refer to either the father or the son). Annotators can split such ambiguous mentions into separate rows and assign more specific names, such as “Alexandre Dumas Sr.” and “Alexandre Dumas Jr.”.

2.2 Relation Extraction

This stage corresponds to the right-hand pane of Figure 1. Its goal is to transform the narrative’s candidate relations \mathcal{E} into a directed graph that captures social, familial, and other types of interactions among the characters, thereby providing a structured representation for subsequent analysis.

Explorer: Edge Candidate Generation. With the character inventory fixed, we run the relation extractor \mathcal{F}_{rel} independently n_e times, each with temperature sampling. Analogous to node extraction, we keep only those triples (x, r, y) that occur

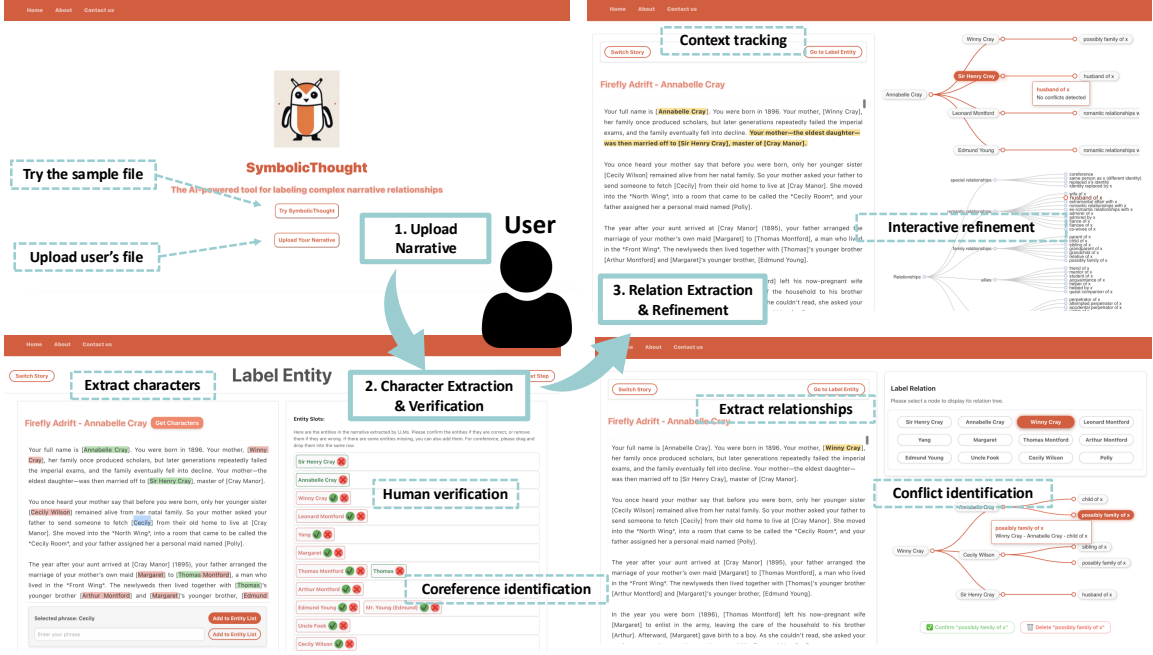


Figure 1: The *SymbolicThought* annotation workflow. Users (1) upload a narrative document, (2) verify or modify automatically extracted character mentions, and (3) confirm, edit, or refine LLM-generated relationships with real-time symbolic consistency checks. The interface provides visual feedback, alerts for logical conflicts, and suggestions for inferable relations, supporting accurate and efficient human-in-the-loop annotation.

in at least τ_e runs, improving recall while suppressing hallucinated relations, where x and y denote character entities and r denotes the relation type:

$$\mathcal{E}_{\text{raw}} = \{(x, r, y) \mid \sum_{i=1}^{n_e} \mathbb{I}((x, r, y) \in \mathcal{F}_{\text{rel}}^{(i)}(\mathcal{D}, \mathcal{V})) \geq \tau_e\}.$$

Symbolic Reasoning LLMs are essentially sequence-based models and often lack a holistic view of the relational graph within a given text. As a result, they frequently miss links when constructing the predicted relation graph (Berglund et al., 2023). Broadly, two types of missing links can occur: (1) **Incomplete relations**, where a unidirectional relationship such as "A to B" is present, but the reverse "B to A" is absent; and (2) **Incomplete structures**, where both "A to B" and "B to A" are missing. Incomplete structures can further be categorised into two subtypes. First, when relationships like "A to B" and "A to C" exist, an inferable relation such as "B to C" (and its reverse "C to B") should be logically deduced, but is missing. Second, although both "A to B" and "B to A" exist, the underlying semantic relations, such as hypernymy or hyponymy, are not identified.

Illustrative examples are provided in Figure 2. To address these issues, we propose auto-completing four categories of inferable relationships, enabling existing triples to entail their logically implied but previously missing counterparts.

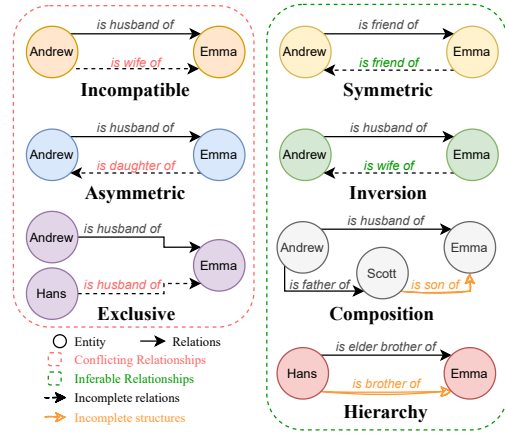


Figure 2: Patterns of interpersonal relationships.

(1) **Symmetry**: if a relationship r_1 holds between x and y , then the same relationship holds between y and x ($r_1(x, y) \implies r_1(y, x)$).

(2) **Inversion**: if a relationship r_1 is present between x and y , then a different relationship r_2 exists between y and x ($r_1(x, y) \implies r_2(y, x)$).

(3) **Composition**: if a relationship r_1 connects x and y , and a relationship r_2 links y and z , then a third relationship r_3 can be inferred between x and z ($r_1(x, y) \cap r_2(y, z) \implies r_3(x, z)$).

(4) **Hierarchy**: Certain relationships can be clas-

sified as subtypes or special cases of more general relationships ($r_1(x, y) \subseteq r_2(x, y)$).

We also identify conflicting relationships, where the presence of one relation precludes another.

(1) Incompatible: The presence of a relationship r_1 between individuals x and y is incompatible with the presence of another relationship r_2 between the same individuals ($r_1(x, y) \implies \neg r_2(x, y)$).

(2) Asymmetric: if a relationship r_1 exists between x and y , then the inverse relationship does not hold between y and x ($r_1(x, y) \implies \neg r_2(y, x)$).

(3) Exclusive: if a specific relationship r_1 exists between individuals x and y , then x cannot have the same relationship with any other individual z ($r_1(x, y) \implies \forall z, \neg r_1(x, z)$).

However, such conflicts often lack clear directionality (i.e., it is unclear which side is incorrect). To resolve this, we retrieve relevant context using a retrieval function $f(r_1, r_2, \mathcal{D})$, where f is implemented via a FAISS-based retriever (Lewis et al., 2020; Douze et al., 2024) that selects supporting evidence from the source text. The retrieved context is then used to construct multiple-choice prompts that help LLMs revise the conflicting relations.

Thus, one of the significant benefits of this annotation and automatic refinement framework is the improved **completeness**, which means that:

Theorem 1 *Let \mathcal{V} denote the set of entities, \mathcal{E} the set of relations, and $\mathcal{G}_o = (\mathcal{V}, \mathcal{E}_o)$ the initial relation graph constructed by the LLM. Given a ground-truth relation graph $\mathcal{G}_g = (\mathcal{V}, \mathcal{E}_g)$, there must exist a finite sequence of operations O_1, O_2, \dots, O_k , where each O_i belongs to the set of operations defined in our symbolic reasoning module, such that: $\mathcal{G}_o \xrightarrow{O_1, O_2, \dots, O_k} \mathcal{G}_g$ always holds.²*

Interactive Human Verification. During annotation, relations suggested by the model are highlighted in yellow. Annotators can confirm these suggestions, turning them green. Relations added manually by annotators are also marked in green by default. For each newly added relation, we automatically perform symbolic reasoning to infer additional triples; these auto-completed relations are again marked in yellow, pending verification. If any conflicting relationships are detected during annotation, they are highlighted in red, with hoverable tooltips showing which existing relation they

conflict with. Additionally, when a user selects a relation to confirm or modify, we highlight the corresponding supporting evidence $f(r_1, r_2, \mathcal{D})$ in the original text to assist their decision.

3 The Relationship Logic Dataset

LLMs on Logical Reasoning We evaluated the performance of LLMs in inferring logical connections among character relationships. The **Add** task focuses on the open-ended inference of new relationships from existing ones, where multiple correct answers exist and the label space is large. We evaluate model performance using *precision*, *recall*, and *F1-score*. The **Remove** task, on the other hand, determines whether a pair of relationships conflict with each other, framed as a three-way classification (*Yes*, *No*, *Unsure*); for this, we report *accuracy* and *F1 score*. Table 1 reveals that even state-of-the-art models struggle to achieve satisfactory results, highlighting that LLMs often lack a grounded understanding of logical relationships and are limited by training data biases. This underperformance underscores the need for more reliable annotation strategies and robust human-in-the-loop frameworks to support consistent and contextually appropriate reasoning in complex narrative texts.

Model	Add			Remove	
	Precision	Recall	F1	Accuracy	F1
GPT-4.1	71.9	55.0	62.3	57.3	42.8
GPT-4o-mini	63.7	48.8	55.3	30.6	36.7
Qwen2.5-32B-Ins	64.7	43.8	52.2	31.6	34.7

Table 1: Performance comparison of LLMs across logical relationship tasks.

Expert Annotation of Logical Relationships To assess LLM performance in logical reasoning over character relationships, two trained experts annotated conflicts and antisymmetric interactions across 160 character pairs, resulting in 51k derived logical relations. Inter-annotator agreement, measured by Cohen’s Kappa (Cohen, 1960), was 0.832, indicating substantial consistency. All disagreements were resolved through discussion.

Corpus Collection and Relation Extraction To facilitate the evaluation of model understanding across diverse textual domains, we curated a corpus of 19 narrative texts spanning four categories: historical accounts, news reports, biographical narratives, and fictional stories. From these, we extracted 1,398 relation triplets among characters (see

²The proof is provided in Appendix C.

Table 5 for details).

Usability Assessment by Non-expert Annotators

To further evaluate the annotation system, we recruited 10 non-expert annotators to assess its usability and user experience. Detailed annotation instructions and additional statistics are provided in Appendix A.

4 Evaluation

4.1 Experimental Setup

We conducted experiments with OpenAI’s GPT-4.1 (gpt-4.1-2025-04-14), GPT-4o-mini (gpt-4o-mini-2024-07-18), and Qwen2.5-32B-Instruct³. For retrieval-augmented generation (RAG), we built the vector library using the text-embedding-3-small model via the OpenAI API⁴. We compare *SymbolicThought* with the following baselines: (1) **Direct Prompting**, where the model is directly prompted with instructions; (2) **Self-Consistency** (Wang et al., 2023), which samples multiple outputs and selects the final answer via majority voting; (3) **Self-Reflection** (Wei et al., 2022), where the model generates initial predictions and then refines them through self-critique; (4) **Human Annotation**, representing fully manual labeling.

4.2 Experimental Results

Character Extraction As shown in Table 2, all models benefit from the self-consistency, we adopt for character extraction, with consistent F1-score gains over direct prompting.

Model	Method	Precision	Recall	F1
GPT-4.1	Prompt	76.9	83.1	79.8
	Self Consistency	77.1	86.7	81.6
GPT-4o-mini	Prompt	82.1	75.8	78.8
	Self Consistency	83.1	77.4	80.2
Qwen2.5-32B-Ins	Prompt	81.0	79.0	80.0
	Self Consistency	83.5	79.4	81.4

Table 2: Performance comparison of models and methods on character extraction.

Relation Extraction. As Table 3 shows, relation extraction from narrative texts remains a highly challenging task for existing LLMs, with all baseline methods yielding relatively low F1 scores, reflecting fragmented reasoning and weak cross-sentence handling. Typical conflicts and further

³<https://huggingface.co/Qwen/Qwen2.5-32B-Instruct>

⁴<https://platform.openai.com/docs/overview>

error analysis are in Appendix D. In contrast, our proposed *SymbolicThought* method, grounded in symbolic reasoning, consistently outperforms all baselines across models.

Human vs. System Performance. For evaluation, we merge system-generated and human annotations to form the final ground truth. As shown in Table 4, *SymbolicThought* achieves substantial improvements over human annotators across all narrative genres, with over 25% higher recall and up to 40% less annotation time. These results confirm the effectiveness of our human-in-the-loop design in streamlining annotation and supporting nuanced social judgments.

4.3 Annotator Feedback Results

To evaluate the usability and effectiveness of *SymbolicThought*, we conducted a user study involving 10 participants. Each participant interacted with the system and completed a structured questionnaire (Figure 3), which covered topics such as interface usability, system performance, feature usefulness, and overall user experience. The complete questionnaire is provided in Appendix B.

As shown in Figure 3, users found the tool highly intuitive and responsive. The majority of users considered the system-generated initial annotations and automatic relation suggestions helpful in reducing their workload. Notably, features such as text highlighting and conflict detection received consistently high usefulness ratings, especially the automatically suggested character relations. While the accuracy of character extraction was generally well received, users reported slightly lower confidence in relation extraction, suggesting opportunities for further refinement. Overall, 9 out of 10 participants indicated they would choose to use the tool in future annotation tasks, confirming the practical value of *SymbolicThought* in supporting efficient and accurate narrative analysis.

4.4 Correlation Analysis: Annotation Complexity vs. Small World Index

Figure 4 plots annotation performance (F1) against the *Small-World Index* (SWI) (Telesford et al., 2017), which is a widely used index from social science research to measure the compactness of a given graph of social relationships. A larger SWI therefore indicates an entity graph that is both locally dense (many closed triads) and globally efficient (short paths), i.e. more “small-world-like”.

Method	GPT-4.1			GPT-4o-mini			Qwen2.5-32B-Ins		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
Prompting	30.9	36.4	33.4	20.6	6.5	9.9	20.9	11.5	14.8
Self-Consistency	29.4	35.0	32.0	27.2	8.1	12.5	23.9	13.1	16.9
Self-Reflection	24.5	26.6	25.5	22.4	7.8	11.5	20.8	12.8	15.8
<i>SymbolicThought</i>	35.1	41.2	37.9	32.2	13.2	18.8	27.7	18.9	22.5

Table 3: Performance comparison of relation extraction using various methods.

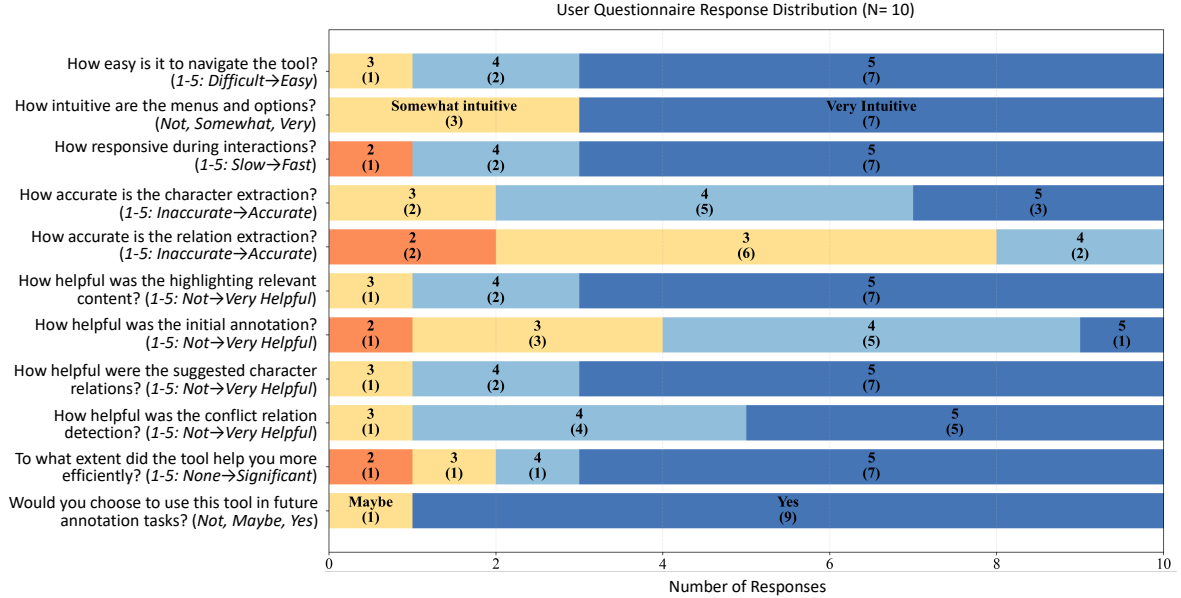


Figure 3: Full Questionnaire Response Distribution. Shades of blue denote more favourable evaluations of the tool, while shades of red indicate dissatisfaction. Grey segments represent neutral factual responses that are not direct indicators of user sentiment toward the tool.




Figure 4: F1 score vs. Small World Index. Each point represents a document, colored by its category.

Cross-category differences. Biography documents cluster in the upper-right portion of the plot, yielding significantly higher F1 scores than Novels and History texts. This suggests that biographies possess more cohesive entity graphs, making them easier to annotate. Novels occupy the middle range, whereas History texts lie predominantly in the lower-left, reflecting their sparser, less structured networks and hence greater annotation

difficulty.

Within-category trend. Across all three categories, F1 increases monotonically with SWI, indicating that richer small-world structure reduces annotation complexity. In practice, SWI thus serves as a reliable proxy for predicting how much human effort a document will require during annotation.

Category	Method	History	Biography	Narrative
Recall (%)	Human	57.3	67.3	63.4
	<i>SymbolicThought</i>	85.6	91.4	89.1
Avg Time (min)	Human	163.4	87.2	102.5
	<i>SymbolicThought</i>	118.3	45.5	74.7

Table 4: Comparison of recall and annotation time between human and our system across different genres.

5 Conclusion

We introduced *SymbolicThought*, a comprehensive human-in-the-loop framework that integrates LLM extraction with symbolic reasoning for more accurate and reliable character relationship annotation.

Our system effectively bridges the gap between model predictions and human expectations, ensuring logical consistency, interpretability, and transparency in the analysis of complex narrative texts.

Limitations

While our proposed method effectively integrates logical reasoning with human-in-the-loop annotation, it still faces several limitations. First, the complexity and uncertainty of relationships across different cultural and historical contexts can challenge the universality of our logical rules. For example, polygamy was considered acceptable in some eras and regions but is not widely accepted today. Our current logic covers the most universal relationships and those found in our dataset, but inevitably falls short of capturing all cultural nuances. Second, the bottleneck of the base models remains a challenge: if the extracted relationship graph contains a high proportion of noisy labels, the logic inference module risks propagating errors, as the inferred relationships are built on incorrect information.

Acknowledgments

This work was supported in part by the UK Engineering and Physical Sciences Research Council (EPSRC) through a Turing AI Fellowship (grant no. EP/V020579/1, EP/V020579/2) and a New Horizons grant (grant no. EP/X019063/1), and KCL's Impact Acceleration Account (grant no. EP/X525571/1). A PhD studentship from the Chinese Scholarship Council funds Qinglin Zhu. The authors also acknowledge the use of the King's Computational Research, Engineering, and Technology Environment (CREATE) at King's College London.

References

- Lukas Berglund, Meg Tong, Max Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. 2023. [The reversal curse: LLMs trained on "a is b" fail to learn "b is a"](#). *CoRR*, abs/2309.12288.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. [The faiss library](#). *CoRR*.
- Qipeng Guo, Xiangkun Hu, Yue Zhang, Xipeng Qiu, and Zheng Zhang. 2023. [Dual cache for long document neural coreference resolution](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, page 1691–1703. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9459–9474.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. [Lost in the middle: How language models use long contexts](#). *CoRR*, abs/2307.03172.
- Abhilasha Sancheti and Rachel Rudinger. 2025. [Tracking evolving relationship between characters in books in the era of large language models](#). In *Proceedings of the The 7th Workshop on Narrative Understanding*, pages 64–82, Albuquerque, New Mexico. Association for Computational Linguistics.
- Qawi K. Telesford, Karen E. Joyce, Satoru Hayasaka, Jonathan H. Burdette, and Paul J. Laurienti. 2017. How small is it? comparing indices of small worldliness. *Network Science*, 5(1):30–44.
- Binghai Wang, Rui Zheng, Lu Chen, Yan Liu, Shihan Dou, Caishuang Huang, Wei Shen, Senjie Jin, Enyu Zhou, Chenyu Shi, Songyang Gao, Nuo Xu, Yuhao Zhou, Xiaoran Fan, Zhiheng Xi, Jun Zhao, Xiao Wang, Tao Ji, Hang Yan, and 8 others. 2024. [Secrets of RLHF in large language models part II: reward modeling](#). *CoRR*, abs/2401.06080.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*. Poster presentation.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Proceedings of the 36th Conference on Neural Information Processing Systems*, New Orleans, United States of America. Curran Associates, Inc.
- Peter West, Ximing Lu, Nouha Dziri, Faeze Brahman, Linjie Li, Jena D. Hwang, Liwei Jiang, Jillian Fisher, Abhilasha Ravichander, Khyathi Chandu, Benjamin Newman, Pang Wei Koh, Allyson Ettinger, and Yejin Choi. 2023. [The generative AI paradox: "what it can create, it may not understand"](#). *CoRR*, abs/2311.00059.

Ruoxi Xu, Yingfei Sun, Mengjie Ren, Shiguang Guo, Ruotong Pan, Hongyu Lin, Le Sun, and Xianpei Han. 2024. *AI for social science and social science of AI: A survey*. *Inf. Process. Manag.*, 61(2):103665.

Xinfeng Yuan, Siyu Yuan, Yuhan Cui, Tianhe Lin, Xintao Wang, Rui Xu, Jiangjie Chen, and Deqing Yang. 2024. *Evaluating character understanding of large language models via character profiling from fictional works*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8015–8036, Miami, Florida, USA. Association for Computational Linguistics.

Runcong Zhao, Qinglin Zhu, Hainiu Xu, Jiazheng Li, Yuxiang Zhou, Yulan He, and Lin Gui. 2024. *Large language models fall short: Understanding complex relationships in detective narratives*. *CoRR*, abs/2402.11051.

A Annotation Details

For *Symmetry & Inversion* relationships, the label space encompassed the complete set of relationships, *R*. For *Asymmetric* and *Incompatible* relationships, the label space consisted of {"Yes", "No", "Unsure"}. The human annotators used the simpler label space {"Yes", "No"}. Because some relationships were challenging to classify, we labelled all disagreements between annotators as "Unsure". For *Composition* relationships, the label space spans all 160 predefined character relationship categories.

In addition to the expert-labeled logical relationships, we also recruited 10 non-expert annotators to test the usability of our annotation system. These annotators were given standardised training materials, identical to the video-based tutorial accompanying this paper. This setup allows us to evaluate the accessibility and efficiency of the system when used by general users rather than domain experts. Each expert annotator completed approximately 40 hours of work and was compensated at a rate of \$31.92 per hour.

Dataset	Historical	Biography	Narrative
#Narratives	4	5	10
Avg. #Character	16.4	24.2	9.0
#Relationships	85.0	130.0	48.7
Avg. #Token	1246.9	737.1	3467.3

Table 5: Dataset Statistics.

B User Questionnaire

To better understand user needs and evaluate the usability and effectiveness of **SymbolicThought**, we

designed a structured questionnaire targeting key aspects of the user experience. The survey covers four major dimensions: interface design, system performance, feature usefulness, and overall feedback. The full questionnaire is provided below.

Section 1: Interface Design

- Q1.** How easy is it to navigate the tool?
(1 = Very Difficult, 5 = Very Easy)
- Q2.** How intuitive are the menus and options?
- Very intuitive
- Somewhat intuitive
- Not intuitive

Section 2: System Performance

- Q3.** How responsive is the tool during interactions?
(1 = Very Slow, 5 = Very Fast)
- Q4.** How accurate is the character extraction provided by the tool?
(1 = Very Inaccurate, 5 = Very Accurate)
- Q5.** How accurate is the relation extraction provided by the tool?
(1 = Very Inaccurate, 5 = Very Accurate)
- Q6.** What is the maximum acceptable wait time for character and relation extraction?
- Less than 1 minute
- Less than 5 minutes
- Depends on the length of the narrative
- Less than the time it would take me to do it manually
- As long as the quality is high

Section 3: Feature Usefulness

- Q7.** How helpful was the system-generated initial annotation (i.e., the AI-suggested pre-filled characters and relation graph) in reducing your workload?
(1 = Not Helpful, 5 = Very Helpful)
- Q8.** How helpful was the text highlighting of relevant content?
(1 = Not Helpful, 5 = Very Helpful)
- Q9.** How helpful were the automatically suggested character relations?
(1 = Not Helpful, 5 = Very Helpful)
- Q10.** How helpful was the conflict relation detection/reminder?
(1 = Not Helpful, 5 = Very Helpful)

Section 4: Overall Feedback

- Q11.** To what extent did the tool help you complete the annotation task more efficiently?
(1 = Not at All, 5 = Significantly)
- Q12.** Would you choose to use this tool in future annotation tasks?
- Yes
- Maybe

C Completeness of operations

To prove **Theorem 1**, we first consider a scenario with the simplest assumption.

Lemma 1 *Theorem 1 is true when there is only one type of relation defined in \mathcal{E} .*

Proof. When there is only one type of relation defined in \mathcal{E} , the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ can be represented as a Boolean adjacency matrix. Specifically, let x_i and x_j be the i -th and j -th entities, respectively. If there exists a relation between them, then the corresponding matrix entry is $\mathcal{G}_{ij} = 1$; otherwise, it is 0.

Both \mathcal{G}_o (the output graph from the LLM) and \mathcal{G}_g (the ground-truth graph) are Boolean matrices. By assumption, \mathcal{G}_o contains a subset of the non-zero entries in \mathcal{G}_g . Therefore, there exists a Boolean matrix \mathcal{G}_a such that:

$$\mathcal{G}_o + \mathcal{G}_a = \mathcal{G}_g$$

where the addition is element-wise Boolean OR.

Since the number of entities is finite, the matrices \mathcal{G}_o , \mathcal{G}_a , and \mathcal{G}_g each contain a finite number of non-zero (i.e., 1) entries.

Under the assumption of a single relation type, the reasoning operations are limited to two forms:

- **T1 (Symmetry):** If $\mathcal{G}_{ij} = 1$, then infer $\mathcal{G}_{ji} = 1$.
- **T2 (Transitivity):** If $\mathcal{G}_{ij} = 1$ and $\mathcal{G}_{jk} = 1$, then infer $\mathcal{G}_{ik} = 1$.

Each operation corresponds to the addition of a single 1-entry in the matrix. Thus, we can construct a finite sequence of such operations to transform \mathcal{G}_o into \mathcal{G}_g in the following steps:

- **Symmetry Completion:** For each pair (i, j) where $\mathcal{G}_{ij} = 1$ appears in \mathcal{G}_o but $\mathcal{G}_{ji} = 1$ only exists in \mathcal{G}_g , we apply a sequence of **T1** operations. Let the number of such pairs be s , and denote the operations as $\sum_{i=1}^s O_i$.
- **Transitive Completion:** For pairs (i, j) such that $\mathcal{G}_{ij} = 1$ appears in \mathcal{G}_g but not in \mathcal{G}_o , and where $i < j$, we apply **T2** operations to infer these missing links. Let the number of such inferences be t , and denote the operations as $\sum_{i=s+1}^{s+t} O_i$.

- **Symmetry for New Edges:** Some of the newly added relations in Step 2 may require symmetric counterparts to satisfy \mathcal{G}_g . We again apply **T1** operations for these cases. Let the number of these be $k - (s + t)$, and denote them as $\sum_{i=s+t+1}^k O_i$.

By composing these operations, we construct a finite sequence O_1, O_2, \dots, O_k such that:

$$\mathcal{G}_o \xrightarrow{O_1, O_2, \dots, O_k} \mathcal{G}_g$$

■

Then, we can extend the **Lemma 1** in to a more complex scenario which contains two types of relations.

Corollary 1 *Theorem 1 is true when there are two types of relations defined in \mathcal{E} .*

Proof. When two relation types are defined in \mathcal{E} , the relation graph can be represented by two corresponding Boolean matrices, each encoding one type of relation. Extending the reasoning in **Lemma 1**, each type of completion (symmetry and transitivity) now applies independently to each relation type. As a result, there are:

- 4 possible cases for symmetry completion (2 relation types \times direction),
- 8 possible cases for transitive completion (combinations across 2 relation types).

In addition, when there are more than one type of relation, we need to consider an additional type of completion, referred to as **copy completion**. This occurs when the presence of a relation of one type implies the presence of a corresponding relation of another type. There are 2 possible cases of copy completion, depending on the direction and pairing of relation types.

Given that the number of entities is finite, and the total number of possible relational triples is bounded, the number of necessary completions remains finite. Thus, it is still possible to construct a finite sequence of operations O_1, O_2, \dots, O_k such that:

$$\mathcal{G}_o \xrightarrow{O_1, O_2, \dots, O_k} \mathcal{G}_g$$

■

Based on the results of **Lemma 1** and **Corollary 1**, the proof of **Theorem 1** follows directly. For multiple relation types, we can first sort the

relations by their frequency of occurrence and then traverse the entire graph. By incrementally editing the relations one by one, we construct a finite sequence of operations holds for all relation types.

D Case Studies

As shown in Figure 5, the initial relation graph extracted by LLMs often contains conflicting information. These include simple incompatibilities (e.g., “child of” vs. “father of”) highlighted in yellow, as well as more subtle exclusive contradictions (e.g., multiple conflicting “husband of” links) highlighted in red. By defining logical constraints over relation types, such as mutual exclusivity or inverse relation, we can automatically identify these inconsistencies in the graph.

To resolve such conflicts, we retrieve supporting evidence from the original narrative using a RAG module. For example, in the top right of Figure 5, a logical conflict is detected between two claims: “Scott is the child of Andrew Palowski” vs. “Scott is the father of Andrew Palowski.” Relevant context is retrieved from the text and presented to the LLM as part of a structured single-choice question. The LLM then selects the logically consistent option based on grounded evidence.

Similarly, in the lower part of the figure, we show an example of exclusive relationship refinement involving two candidates for the husband of Andrew Palowski’s wife. The tool retrieves narrative evidence from multiple perspectives and uses contrastive prompting to resolve the ambiguity. In both examples, the system leverages the symbolic logic layer to surface conflicts and guide the LLM’s reasoning, while maintaining human interpretability through structured prompts and explicit evidence grounding.

This refinement pipeline allows *Symbolic Thought* to iteratively improve the coherence and correctness of the extracted relation graph, addressing one of the key limitations of LLMs, namely, inconsistent or contradictory relational inference in complex narrative contexts.

D.1 Error Analysis

Below, we categorise and analyse the common failure modes observed during our experiments.

What kind of relationships do LLMs miss?

LLMs often miss relationships that require some inference, such as “*sister-in-law*”, “*brother-in-law*”, “*father-in-law*”, etc. These relationships are often

not explicitly stated in the original text. For example, A is B’s wife, and B is C’s child, so A should be C’s daughter-in-law. In such cases, LLMs often fail to successfully capture this information.

What kind of hallucinations do LLMs generate?

There are differences in how humans and machines understand relationships. For example, consider the following text: “*When you were in dire straits, you often confided in Father Tom, who visited Brighton Hospital. You told him your misfortunes, your drinking problem. Tom tried his best to comfort you, and on his advice, you participated in a treatment program led by a psychiatrist...*”. In this context, human annotators believe that the priest helped Andrew, so he is a helper and a saviour. However, LLMs perceive him as an information receiver from Andrew. While this interpretation is not entirely incorrect, it is a more superficial relationship. Such differences in understanding lead to many noisy labels generated by LLMs.

What kind of logic mistakes do LLMs make?

One issue is that LLMs often struggle to identify the direction of relationships. We tried providing few-shot examples and instructions, but these approaches did not yield significant improvements. Another issue is that when we identify a conflicting relationship and ask the model to reflect on which is correct, it often answers both incorrectly, displaying strong biases. To address this, we adopt a multiple-choice format that forces the model to choose the most likely option.



Figure 5: A demonstration of our relationship refinement pipeline. Starting from rough predictions (top left figure), we leverage the logic relationships to identify erroneous relationship edges. Depending on the type of logical conflict, we support LLMs with context retrieved via RAG and a multiple-choice prompt to refine the relationships.