IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 2 September 2025, accepted 22 October 2025, date of publication 28 October 2025, date of current version 6 November 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3626479

RESEARCH ARTICLE

DYNA-SKILL: Dynamic Self-Prompting Knowledge Graphs for Improving Logical Reasoning in Language Models

JONGWON RYU, (Student Member, IEEE), MINGI KIM, (Student Member, IEEE), AND JUNYEONG KIM, (Member, IEEE)
Department of Artificial Intelligence, Chung-Ang University, Seoul 06974, Republic of Korea

Corresponding author: Junyeong Kim (junyeongkim@cau.ac.kr)

This work was partly supported by Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea Government (MSIT) (No. RS-2022-II220184, Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics) and partly supported by the Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea Government (MSIT) [RS-2021-II211341, Artificial Intelligence Graduate School Program (Chung-Ang University)].

ABSTRACT Despite significant advances in large language models, many reasoning datasets are still built from a fixed set of predefined relations, manually curated types such as cause, effect, and intent found in knowledge graph datasets such as ATOMIC and COMET. While these predefined relations provide essential structure, the fixed schema limits relational coverage and adaptability to novel contexts. We present DYNA-SKILL, a dual-triple knowledge graph framework that preserves 35 predefined relations consolidated and refined from existing commonsense knowledge graph datasets while augmenting them with 133 additional schema-free dynamic relations generated via a self-prompting mechanism. Each instance consists of two linked triples (Head–Predefined Relation–Tail) and (Tail–Dynamic Relation–Additional Tail) used as independent training samples while retaining linkages for extended reasoning paths. Across reasoning-intensive benchmarks, including CommonsenseQA, RiddleSense, and ARC Challenge, the Hybrid configuration, which combines predefined and dynamically generated relations, achieves performance comparable to or slightly higher than Predefined-only settings and yields up to 3.2% higher accuracy than baseline BERT models. By expanding the relation set from 35 predefined types to a total of 168 relations, DYNA-SKILL enriches relational diversity and improves multi-step logical reasoning, which can enhance performance in real-world scenarios such as complex question answering, multi-document analysis, and causal reasoning, where accurate and adaptable inference is critical.

INDEX TERMS Knowledge graphs, logical reasoning, language models, dynamic relations, self-prompting.

I. INTRODUCTION

In recent years, large language models (LLMs) have demonstrated remarkable capabilities across a wide range of natural language processing tasks, including question answering, summarization, and commonsense reasoning [1]. Despite these advances, LLMs continue to struggle with complex, multi-step logical reasoning, particularly in open-domain and contextually rich scenarios [2]. This limitation is partly due to their reliance on implicit knowledge learned during pretraining, without explicit relational structures that facilitate structured inference.

A large proportion of reasoning evaluation datasets are still constructed from a fixed set of manually curated relation types, such as cause, effect, and intent, found in commonsense knowledge graph datasets like ATOMIC [3], [4] and expanded using models such as COMET [5]. While these predefined relations provide essential structure,

The associate editor coordinating the review of this manuscript and approving it for publication was Zijian Zhang.

© 2025 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License.
For more information, see https://creativecommons.org/licenses/by/4.0/

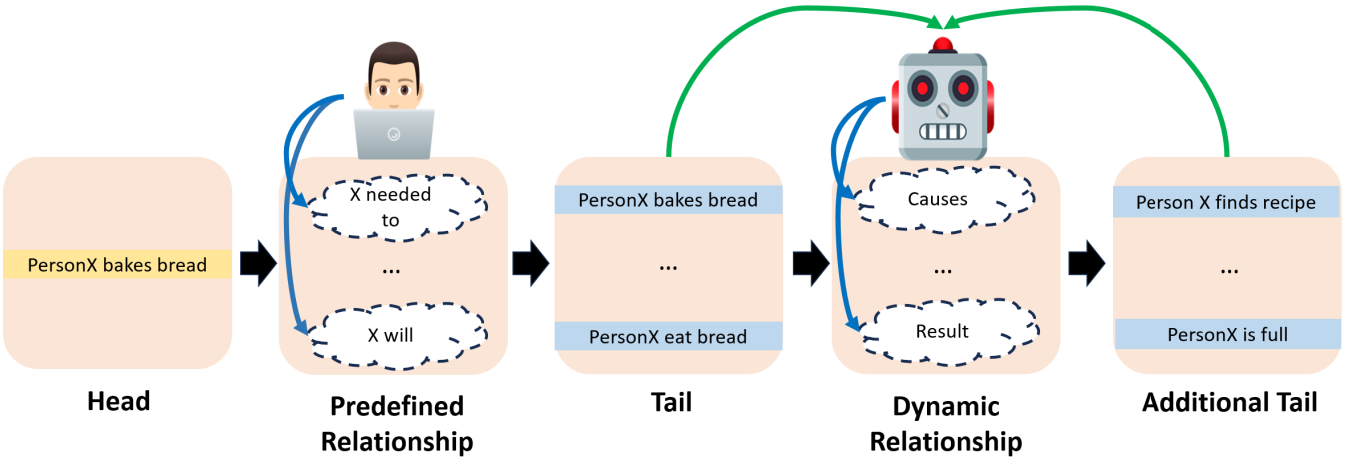


FIGURE 1. Overview of the self prompting graph based knowledge dataset generation.

the fixed schema inherently limits relational coverage and adaptability to novel or context-specific connections. As a result, current reasoning datasets cannot fully support the diverse and dynamic relational patterns required for robust, multi-step inference in real-world applications.

To address these limitations, we propose **DYNA-SKILL**, a Self-Prompting-based approach for automatically generating a graph-structured knowledge dataset that integrates both predefined and dynamically generated relations. As illustrated in Figure 1, our method constructs a dual-triple representation: (Head–Predefined Relation–Tail) and (Tail–Dynamic Relation–Additional Tail). First, Tails are generated using 35 predefined relations, consolidated and refined from existing commonsense knowledge graph datasets. To extend reasoning depth, an Additional Tail is generated based on the Tail, introducing new but logically coherent knowledge. Finally, a Dynamic Relation is inferred between the Tail and Additional Tail, enabling the discovery of 133 schema-free relation types beyond manually curated templates. For example, given the Tail "PersonX bakes bread," our approach may generate the Additional Tail "PersonX finds a recipe" and infer the Dynamic Relation "Causes."

This Self-Prompting-driven process enables LLMs to learn diverse and flexible relational structures, facilitating multi-step inference and contextually adaptive reasoning across various domains. The dual-triple structure serves as a foundation for enhanced logical reasoning, capturing both explicit and implicit connections that conventional knowledge graphs often miss.

We evaluate DYNA-SKILL on five well-established reasoning benchmarks: ARC Challenge [6], CommonsenseQA [7], HellaSwag [8], QASC [9], and RiddleSense [10]. Additionally, we compare against a control dataset (CC News) to isolate the specific contribution of our reasoning-focused dataset beyond general language understanding. Our results show that models fine-tuned on DYNA-SKILL consistently outperform both baseline and control models, particularly in tasks that require multi-step inference.

The main contributions of this work are as follows:

- We introduce a method for automatically generating a graph-based knowledge dataset that integrates 35 predefined and 133 dynamically generated relations, substantially increasing the adaptability and coverage of the knowledge base.
- We develop a dual-triple structure (Head–Relation–Tail and Tail–Dynamic Relation–Additional Tail) that supports multi-step inference and captures a broader range of logical relationships beyond existing commonsense graphs.
- Through experiments on multiple reasoning benchmarks, we demonstrate that DYNA-SKILL significantly enhances LLMs' logical reasoning performance, outperforming both baseline and control models, thereby validating the effectiveness of our approach.

II. RELATED WORK

Our work connects two previously distinct lines of research: (1) reasoning-specific datasets in the form of structured commonsense knowledge graphs, and (2) dynamic relation generation methods such as self-prompting. While knowledge graphs like ATOMIC and COMET provide explicit relational structures, they are restricted by fixed relation schemas. Conversely, dynamic generation methods offer adaptability but lack integration with structured, reasoning-specific graph formats. DYNA-SKILL combines these strengths by embedding schema-free dynamic relations within a predefined relational framework, creating a flexible yet structured resource for logical reasoning.

A. REASONING-SPECIFIC KNOWLEDGE GRAPH DATASETS

1) ATOMIC

ATOMIC [3] is one of the first large-scale commonsense knowledge graphs tailored for "if–then" reasoning. It captures human-centered scenarios through categories such as intentions, reactions, and effects. Its manually curated triples

ensure high quality, but the fixed set of relation types limits coverage and adaptability to novel contexts.

2) COMET

COMET [4], [5] extends ATOMIC by using transformer-based models to populate predefined relational templates derived from ATOMIC and ConceptNet. Although this automates triple generation, COMET remains bound to its original set of predefined relations, preventing adaptation to unseen relation types.

3) OTHER COMMONSENSE GRAPHS

ConceptNet [11] and other large-scale resources such as Freebase [12], DBpedia [13], and YAGO [14] cover a wide range of domains, but their relation inventories are static and schema-bound, which constrains their use for tasks requiring dynamically evolving logical connections.

B. DYNAMIC RELATION GENERATION METHODOLOGIES

Self-prompting approaches, such as [15], are not designed to construct reasoning-specific graph datasets. Instead, they dynamically generate contextually relevant prompts and answers in multi-step open-domain QA. While effective for adaptive knowledge acquisition, these methods typically operate without an underlying structured graph, limiting their ability to produce explicit multi-step relational chains for reasoning.

III. METHOD

In this study, we present DYNA-SKILL, a graph-based knowledge dataset designed to enhance the logical reasoning capabilities of language models. Using a Self-Prompting approach [15] with the GPT-4-turbo API [16], we automatically construct dual-triple knowledge structures in the form of (Head–Predefined Relation–Tail) and (Tail–Dynamic Relation–Additional Tail). Each component, Head, Tail, Dynamic Relation, and Additional Tail, is generated to ensure contextual relevance and relational diversity. Figure 2 illustrates the overall pipeline, from data generation to fine-tuning and evaluation. The following subsections detail each stage of the methodology, and illustrative examples of such dual-triple structures are provided in Table 1.

Algorithm 1 DYNA-SKILL Dataset Generation Process

Input: Predefined relation set R_{pre} (35 types), Head category set C, Large Language Model M
Output: Dual-triple dataset D
Initialize $D \leftarrow \emptyset$
for all category $c \in C$ **do**
  Head Generation: Select category c and generate Head h using M with a category-specific prompt.
  Relation Selection: Choose a predefined relation $r_{pre} \in R_{pre}$ that matches the semantic type of h.
  Tail Generation: Generate Tail $t \leftarrow M(\text{prompt}(h, r_{pre}))$ using a relation-specific template (e.g., "What is the typical use of *Head*?" for *ObjectUse*).
  Additional Tail Generation: Generate Additional Tail $t_{add} \leftarrow M(\text{prompt}(t))$ to extend reasoning depth.
  Dynamic Relation Inference: Infer Dynamic Relation $r_{dyn} \leftarrow M(\text{relation-prompt}(t, t_{add}))$ using a schema-free relation prompt (e.g., "What is the relationship between *Tail* and *Additional Tail*?").
  Append both triples (h, r_{pre}, t) and (t, r_{dyn}, t_{add}) to D.
end for
Text Conversion: For each triple in D, convert to a natural language sentence using predefined mapping rules.
Data Storage: Save the converted sentences to a plain text file for fine-tuning.
return D

1) **Head-to-Tail Generation**
Head Definition: We define *Head* entities across diverse categories to represent the main subjects of logical reasoning events. The categories cover a broad range of commonsense scenarios, including:
   - **Social Interaction** (e.g., education, household, relationship management)
   - **Physical Entities** (e.g., tools, vehicles, appliances)
   - **Event-Centered** (e.g., festivals, weddings, public gatherings)
   - **Causal Relations** (e.g., economic events, technological failures, climate events)

   Additional categories include Causal Chain, Temporal Relations, Duration, Frequency, Direction and Movement, Conditional Relations, Necessary and Sufficient Conditions, Hierarchical Relations, Part-Whole Relations, and Quantitative Relations. Each category captures distinct logical structures and interactions, ensuring coverage of diverse reasoning contexts.
   Relation Definition: Each *Head* category is associated with predefined relations that guide the generation of Tail elements and ensure consistency across the dataset. These relations include context-specific types tailored to each category. Drawing on insights from prior works such as ATOMIC and COMET, we expand the range of predefined relations to build a richer and more varied relational structure.
   - **Social-Interaction Relations** Examples: *xIntent* (intention behind an action), *xNeed* (prerequisites for an action), *oEffect* (impact on others) These relations capture interpersonal and motivational aspects, enabling reasoning about complex social dynamics.
   - **Physical-Entity Relations** Examples: *ObjectUse* (typical use of an object), *AtLocation* (where an object is typically found), *CapableOf* (actions an object can perform) These describe functional

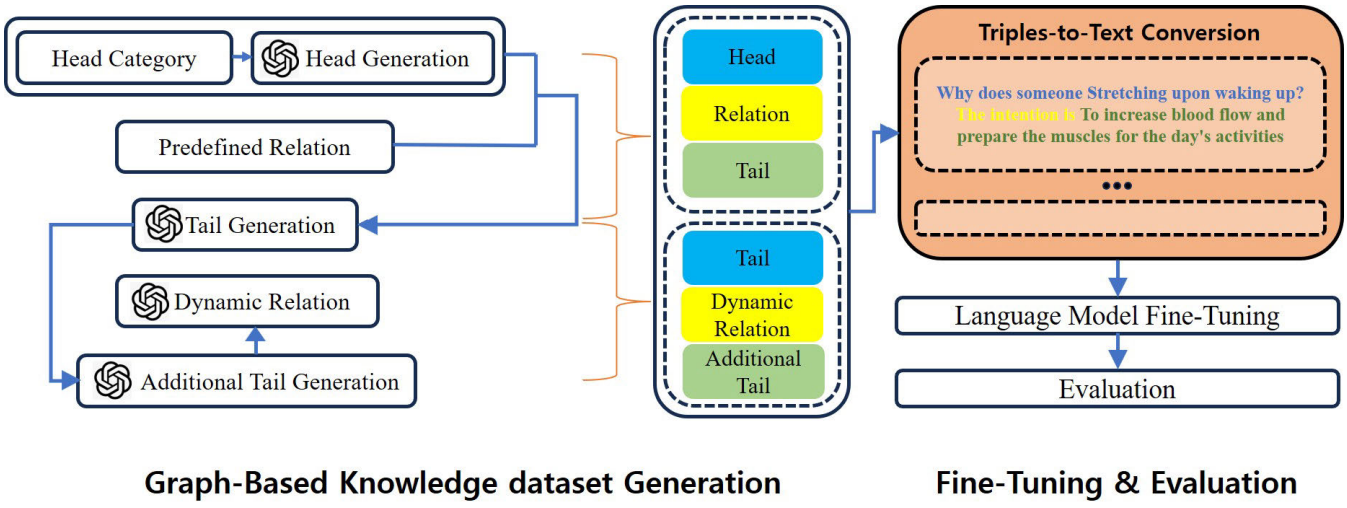


FIGURE 2. Overview of the DYNA-SKILL framework. The dataset generation phase creates dual-triple knowledge structures by combining predefined and dynamically generated relations. These are converted into natural language for fine-tuning, and models are evaluated on reasoning benchmarks.

and situational properties essential for practical reasoning.

- **Event-Centered Relations** Examples: *IsAfter* (what happens after an event), *HasSubEvent* (sub-events of a main event), *Causes* (what leads to an event) These support temporal and causal reasoning beyond fixed templates.
- **Causal Relations** Examples: *Cause and Effect*, *Causal Chain* These describe outcome dependencies and multi-step cause–effect sequences.
- **Other Categories** Examples: *Temporal Sequence* (Temporal Relations), *If-Then Statements* (Conditional Relations), *Part-Whole Relations* (compositional structures), *Quantities and Measures* (Quantitative Relations) These model temporal dependencies, conditional logic, and hierarchical structures.

Each relation is paired with a specific prompt to guide Tail generation. For example, an *xIntent* relation for a social action Head may use the prompt: *"What is the possible intention behind this action?"*

By extending relation types beyond those in existing commonsense graphs, we provide a versatile framework that supports richer logical connections, including cause–effect, hierarchical, and conditional dependencies.

2) **Tail-to-Additional Tail and Dynamic Relation Generation**

Additional Tail Generation: To extend the initial Tail, we apply a Self-Prompting approach to generate an *Additional Tail* that is contextually related to the existing Tail. This step deepens logical connections by prompting the model with targeted questions about further related actions, events, or consequences.

Dynamic Relation Generation: We then determine the relationship between the Tail and the Additional Tail by asking the model: *"What kind of relationship does 'additional tail' have with 'tail'?"* This enables the automatic creation of previously undefined, schema-free relations, thereby enhancing flexibility and incorporating novel, context-specific connections into the dataset.

3) **Dual-Triple Structure: (Head – Relation – Tail) and (Tail – Dynamic Relation – Additional Tail)**

Triple Separation: Each data point is structured as two distinct triples: (Head, Relation, Tail) and (Tail, Dynamic Relation, Additional Tail) This dual-triple structure enables multi-step reasoning by connecting events in layered logical relationships.

Multi-Layered Logical Representation: The dual-triple format allows the representation of complex, multi-step relationships that go beyond simple fact-based connections, enabling the language model to learn deeper logical reasoning capabilities.

4) **Text Conversion of Triples for Language Model Fine-Tuning**

Triple-to-Text Conversion: After generating the (Head, Relation, Tail) triples, we convert each triple into a natural language sentence using a function designed to adapt each relation type into a specific sentence structure. For example, a triple such as:

Head: "PersonX makes coffee", Relation: "xIntent", Tail: "to help"

is converted into:

"Why does someone make coffee? The intention is to help."

Conversion Process: A hybrid-relation function processes each triple according to its relation type, producing readable sentences. This ensures that each triple is expressed as a coherent and contextually relevant sentence that is easy for the language model to interpret.

TABLE 1. Examples of Dual-Triple Structures with Natural Language Conversion based on the conversion rules in the text processing script. Each example consists of two linked triples: (Head–Predefined Relation–Tail) followed by (Tail–Dynamic Relation–Additional Tail), illustrating how predefined and dynamically generated relations connect to form extended reasoning paths.

Domain	Triple 1 (Predefined Relation)	Triple 2 (Dynamic Relation)
Sports Action	*(Spiking – xNeed – A set or a pass is needed before performing a spike in volleyball.)* **Sentence:** What does someone need before Spiking? A set or a pass is needed before performing a spike in volleyball.	*(A set or a pass is needed before performing a spike in volleyball. – Preparatory – Pulling back the bowstring before releasing an arrow in archery.)* **Sentence:** A set or a pass is needed before performing a spike in volleyball and Pulling back the bowstring before releasing an arrow in archery are connected by Preparatory.
Safety Action	*(Shooting – xNeed – Safety training and proper authorization or permits.)* **Sentence:** What does someone need before Shooting? The prerequisite is Safety training and proper authorization or permits.	*(Safety training and proper authorization or permits. – Complementary – Conducting regular safety audits and inspections.)* **Sentence:** What complements Safety training and proper authorization or permits? It is complemented by Conducting regular safety audits and inspections.
Physical Object	*(Cans – CapableOf – Cans are capable of storing and preserving food or liquids.)* **Sentence:** What is Cans capable of? Cans are capable of storing and preserving food or liquids.	*(Cans are capable of storing and preserving food or liquids. – Functional – Dehydrating fruits and vegetables.)* **Sentence:** Cans are capable of storing and preserving food or liquids and Dehydrating fruits and vegetables are connected by Functional.
Leisure Activity	*(Snorkeling or scuba diving – oWant – Others might want to try snorkeling or scuba diving themselves.)* **Sentence:** What do others want after Snorkeling or scuba diving? Others might want to try snorkeling or scuba diving themselves.	*(Others might want to try snorkeling or scuba diving themselves. – Alternative – Sailing)* **Sentence:** Others might want to try snorkeling or scuba diving themselves and Sailing are connected by Alternative.
Scientific Material	*(Polyester – MadeUpOf – Polyester is a synthetic polymer made primarily from petroleum-derived ethylene glycol and terephthalic acid.)* **Sentence:** What is Polyester made up of? Polyester is a synthetic polymer made primarily from petroleum-derived ethylene glycol and terephthalic acid.	*(Polyester is a synthetic polymer made primarily from petroleum-derived ethylene glycol and terephthalic acid. – Chemical – Synthesis of polycarbonate from bisphenol A and phosgene.)* **Sentence:** Polyester is a synthetic polymer made primarily from petroleum-derived ethylene glycol and terephthalic acid and Synthesis of polycarbonate from bisphenol A and phosgene are connected by Chemical.

Storing and Preparing Data for Fine-Tuning: The converted text data is saved line-by-line in a text file, which serves as the input for fine-tuning. This conversion enables the dataset to be directly utilized in model training, enhancing logical reasoning capabilities through structured, narrative-like training data.

5) **Fine-Tuning Language Models on Converted Text Data**

Fine-Tuning Setup: We fine-tune BERT, RoBERTa, DeBERTa, and DistilBERT models [17], [18], [19], [20] using the converted text data. Each model is trained to enhance its logical reasoning capabilities with our dataset, which provides explicit logical connections.

Comparison with Control Dataset: To evaluate the specific contribution of our dataset to logical reasoning, we compare the performance of models fine-tuned on our hybrid-relation dataset with those fine-tuned on a control dataset (CC News), which is expected to have limited impact on logical reasoning. By observing that models trained on CC News show a smaller improvement in logical reasoning tasks compared to those trained on our dataset, we demonstrate that our dataset effectively enhances reasoning capabilities in a way that general text data cannot. This comparison underscores the value of our graph-structured knowledge in fostering deeper inference abilities.

IV. EXPERIMENTS

A. DATASETS

1) HYBRID-RELATION GRAPH DATASET

Our primary dataset is generated via the Self-Prompting method described in Section III, combining 35 predefined relations with 133 dynamically generated relations in a dual-triple format. These triples are converted into natural language sentences for model fine-tuning. We use approximately 100,000 sentences for each training run.

2) PREDEFINED-ONLY GRAPH DATASET

A variant of the above dataset containing only the 35 predefined relations, without any dynamically generated relations. The dataset size is matched to the others at approximately 100,000 sentences, allowing a fair comparison to isolate the contribution of dynamic relation generation.

3) CC NEWS DATASET

A large-scale news corpus used as a control dataset for general-domain fine-tuning. Lacking a reasoning-specific structure, it is expected to have limited impact on logical reasoning performance. We randomly sample 100,000 sentences for size parity with the other datasets.

B. MODELS

We evaluate four transformer-based encoder models with different capacities: BERT [17], RoBERTa [18], DeBERTa [19],

and DistilBERT [20]. This selection allows us to measure dataset impact across both large and lightweight architectures.

C. EVALUATION TASKS

To assess improvements in logical reasoning, we employ five established benchmarks: ARC-Challenge [6], CommonsenseQA [7], HellaSwag [8], QASC [9], and RiddleSense [10]. These tasks cover diverse reasoning types, including multiple-choice science questions, commonsense inference, situational plausibility, multi-hop QA, and lateral thinking riddles.

D. FINE-TUNING AND EVALUATION PROCEDURE

For each model–dataset pair, we fine-tune using approximately 100,000 training sentences, keeping dataset sizes consistent to control for size effects. Fine-tuning is conducted with a standard language modeling objective, batch size 32, learning rate 2×10^{-5}, and 3 epochs. Evaluation is performed on the benchmark test sets, with accuracy as the primary metric for all tasks.

E. COMPARISON SETTINGS

We conduct a two-tier comparison:

1) **Baseline vs. Graph Datasets:** Comparing models fine-tuned on each graph dataset against their unfine-tuned baselines.
2) **Hybrid vs. Predefined vs. CC News:** Comparing reasoning gains from dynamic relations (Hybrid), static relations (Predefined), and general-domain fine-tuning (CC News) to determine the specific contribution of dynamic relation generation.

All datasets contain the same number of sentences, ensuring differences are attributable to content rather than size.

V. RESULT

A. PERFORMANCE COMPARISON ACROSS BASELINE AND HYBRID-RELATION FINE-TUNED MODELS

Table 2 compares baseline models with those fine-tuned on the hybrid-relation Graph Dataset across five reasoning benchmarks. Overall, hybrid-relation fine-tuning yields consistent gains over baseline performance, with notable improvements on ARC-Challenge, CommonsenseQA, and RiddleSense. These gains suggest that the dataset's structured, multi-relational design supports more effective multi-step inference and nuanced commonsense reasoning.

While improvements on QASC and HellaSwag are smaller, the results indicate that hybrid-relation fine-tuning still maintains competitive performance, highlighting potential for further enhancement by integrating richer contextual or domain-specific knowledge.

B. COMPARISON BETWEEN HYBRID-RELATION AND CC NEWS FINE-TUNING

Table 2 compares models fine-tuned on the hybrid-relation Graph Dataset with those fine-tuned on the CC News Dataset, isolating the effect of reasoning-specific data. Across most benchmarks, hybrid-relation fine-tuning yields higher scores on tasks such as ARC-Challenge, CommonsenseQA, and RiddleSense, indicating that structured, multi-relational knowledge directly benefits logical inference.

Model-specific trends further support this conclusion. For example, BERT shows clear gains on ARC-Challenge and RiddleSense when trained on the hybrid-relation dataset, while RoBERTa achieves higher accuracy on CommonsenseQA and HellaSwag, suggesting improved situational and commonsense reasoning. DeBERTa exhibits consistent advantages across tasks, with notable improvements in HellaSwag and QASC, reinforcing the dataset's utility for multi-step inference.

These results validate the hypothesis that reasoning-focused datasets offer advantages over general-purpose corpora for logical reasoning. While CC News improves general language understanding, it lacks the explicit relational structures needed to support complex, stepwise reasoning.

C. COMPARISON BETWEEN HYBRID-RELATION AND PREDEFINED-ONLY FINE-TUNING

Table 2 also reports results for a *Predefined-Only* variant of our dataset, containing the same 100,000 samples but restricted to the 35 predefined relations without any dynamically generated ones. This comparison isolates the contribution of dynamic relation generation to reasoning performance.

Overall, the performance gap between the Hybrid-Relation and Predefined-Only settings is modest but consistent across several tasks. For example, BERT shows small gains on ARC-Challenge and RiddleSense with Hybrid-Relation training, while RoBERTa benefits slightly on CommonsenseQA and QASC. DeBERTa and DistilBERT also exhibit minor but positive differences in most benchmarks, suggesting that dynamically generated relations introduce additional contextual variety that can support reasoning beyond the coverage of fixed relations.

Although the improvements are not large in absolute terms, their presence across multiple architectures and tasks indicates that dynamic relations add complementary knowledge that predefined schemas cannot fully capture. These results imply that even small increments in relational diversity can compound over multi-step reasoning chains, leading to more robust inference capabilities.

D. QUALITATIVE ANALYSIS OF DYNAMIC RELATIONS

The hybrid-relation Graph Dataset incorporates a diverse set of dynamically generated relations, adding flexibility to the model's reasoning capabilities. By filtering out relations that appear fewer than ten times, we identified 133 unique dynamic relations, which occur a total of 49,998 times throughout the dataset. The most frequently occurring relation was *Causal*, appearing 24,825 times, but as this is a pre-existing relation, we excluded it from the analysis of novel dynamic relations. Figure 3 shows

TABLE 2. Accuracy (%) of each model on five reasoning benchmarks. Bold indicates the best score and underline the second best within each model type. Δ represents the accuracy difference between Hybrid and Predefined-only settings. Hybrid-relation fine-tuning generally achieves competitive or superior results, suggesting that dynamic relation generation contributes to improved logical reasoning performance.

	ARC-Challenge [6]	Commonsense QA [7]	HellaSwag [8]	QASC [9]	Riddle Sense [10]
BERT [17]	22.61	18.76	24.59	11.12	19.59
BERT_CC_NEWS [17]	24.83	19.33	**24.72**	**13.50**	18.41
BERT_Hybrid [17]	**25.77**	20.64	24.60	11.56	**20.37**
BERT_Predifined [17]	25.71	**20.65**	24.58	11.59	20.33
Δ (Hybrid-Predefined)	+0.06	-0.01	+0.02	-0.03	+0.04
RoBERTa [18]	25.43	19.00	24.69	**13.82**	16.69
RoBERTa_CC_NEWS [18]	**26.88**	21.05	25.19	12.63	17.92
RoBERTa_Hybrid [18]	23.72	**22.52**	**25.44**	11.66	**19.78**
RoBERTa_Predefined [18]	23.70	22.49	25.43	11.61	19.77
Δ (Hybrid-Predefined)	+0.02	+0.03	+0.01	+0.05	+0.01
DeBERTa [19]	23.04	19.08	24.84	11.99	**21.25**
DeBERTa_CC_NEWS [19]	**25.60**	19.66	24.36	11.66	17.53
DeBERTa_Hybrid [19]	25.09	**20.39**	**25.62**	**12.74**	18.51
DeBERTa_Predefined [19]	25.11	20.36	25.61	12.71	18.52
Δ (Hybrid-Predefined)	-0.02	+0.03	+0.01	+0.03	-0.01
DistilBERT [20]	**25.77**	18.84	24.76	12.53	**21.84**
DistilBERT_CC_NEWS [20]	25.34	18.59	25.15	12.42	20.67
DistilBERT_Hybrid [20]	23.55	**19.49**	**25.71**	**13.07**	17.60
DistilBERT_Predefined [20]	23.54	19.44	**25.71**	13.05	17.58
Δ (Hybrid-Predefined)	+0.01	+0.05	0.00	+0.02	+0.02

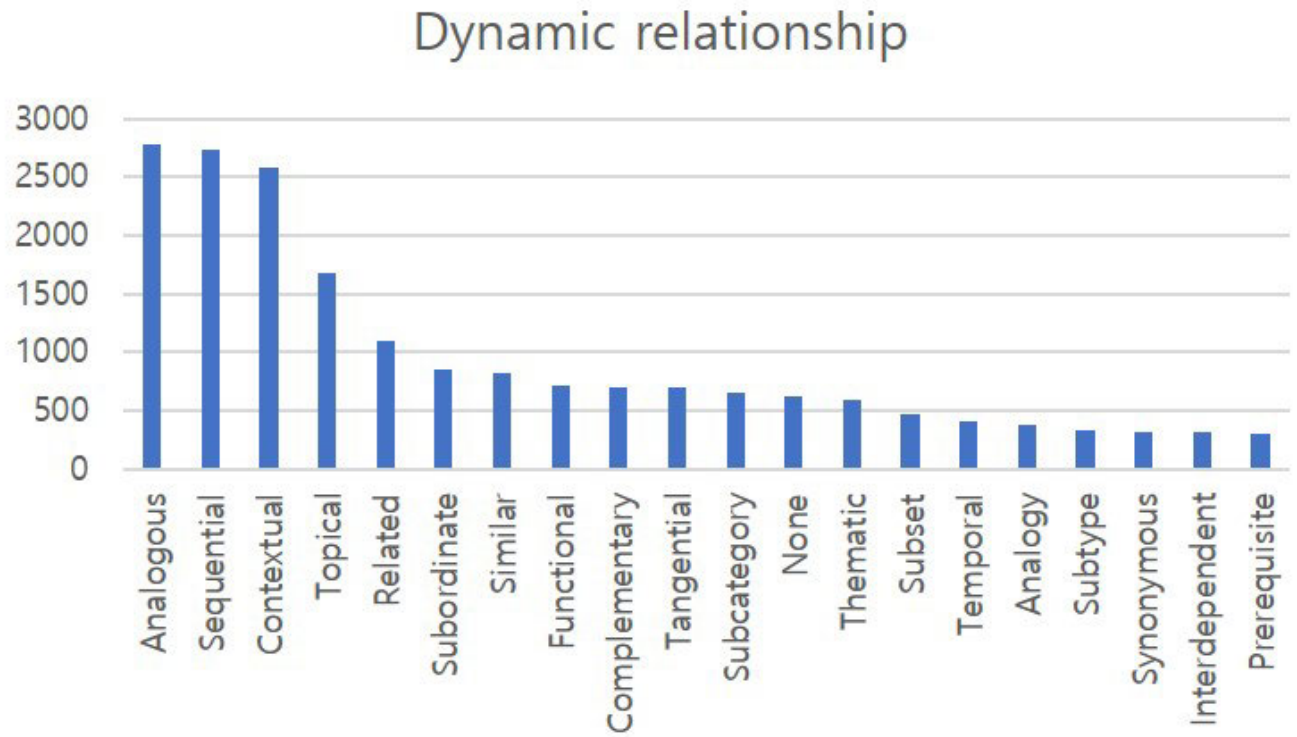


FIGURE 3. Frequency distribution of the top 20 dynamic relation types (excluding the single most common type). Relations such as ***Analogous***, ***Sequential***, and ***Contextual*** occur most frequently, indicating that the self-prompting generation process captures a broad spectrum of context-specific and non-predefined connections. This variety reflects the dataset's ability to extend beyond fixed schemas and enrich multi-step reasoning.

the top 20 dynamic relations ranked from the 2nd to the 21st most frequent, with types like *Analogous*, *Sequential*, *Contextual*, and *Complementary* appearing most frequently. These relations support nuanced, multi-step reasoning by creating contextually rich connections between concepts. These dynamic relations offer models additional relational

context, enabling them to make logical inferences that extend beyond standard, predefined relational structures.

VI. DISCUSSION

Our results show that the hybrid-relation Graph Dataset consistently enhances logical reasoning performance across multiple transformer-based architectures [17], [18], [19], [20], validating the benefit of combining predefined and dynamically generated relations in a graph-structured format. The inclusion of 133 schema-free dynamic relations, in addition to 35 predefined types, enables richer multi-step and causal reasoning than fixed-schema datasets alone.

A. EFFECTIVENESS ACROSS BENCHMARKS

In the first comparison, models fine-tuned on the hybrid-relation dataset outperformed baseline models on reasoning benchmarks [6], [7], [8], [9], [10], [15], particularly CommonsenseQA, RiddleSense, and ARC-Challenge. After filtering low-frequency relations, 133 unique dynamic types remained across 49,998 instances, with frequent categories including *Causal*, *Analogous*, and *Contextual*. These relations provide diverse inference pathways, supporting more flexible reasoning.

B. COMPARISON WITH GENERAL-PURPOSE DATA

In the second comparison, the hybrid-relation dataset generally outperformed CC News on reasoning tasks, confirming that explicit relational structure yields unique benefits. Nonetheless, in QASC and HellaSwag, CC News achieved comparable or slightly higher scores, suggesting that broad-domain knowledge can still aid certain forms of inference. This points to the potential of hybrid training strategies that integrate reasoning-specific and general-purpose data.

C. MODEL CAPACITY CONSIDERATIONS

Larger models such as RoBERTa and DeBERTa benefited more from the structured dataset than smaller models like DistilBERT, indicating that model capacity influences the ability to leverage complex relational structures. For resource-limited settings, simplified or distilled variants of the dataset may be necessary to deliver similar benefits.

D. QUALITATIVE INSIGHTS AND CHALLENGES

Dynamic relations extend coverage beyond fixed schemas and capture nuanced, context-specific links absent in traditional commonsense graphs. However, automatic generation can produce inconsistencies or overly broad labels. Refining prompt design and incorporating automated validation mechanisms could improve precision and alignment with task requirements.

E. LIMITATIONS AND FUTURE WORK

This study used 100K instances for each dataset. While effective, this scale may not fully capture the diversity of logical relations needed for more complex tasks. Future work will expand the dataset to 300K instances and conduct balanced comparisons against equivalently scaled CC News data to assess the interaction between dataset size and reasoning performance. Moreover, the scalability of the approach to significantly larger and noisier real-world datasets, especially those with highly heterogeneous relation types, remains an open challenge that warrants further investigation.

Our evaluation focused on reasoning benchmarks; transferability to other domains, such as fact verification or knowledge retrieval, remains unexplored. Exploring cross-domain applicability, along with model–dataset co-design strategies for smaller architectures, represents an important direction. Finally, while self-prompting allows flexible generation of dynamic relations, ensuring their logical validity remains an open challenge that warrants targeted verification techniques.

VII. CONCLUSION

In this study, we presented the hybrid-relation Graph Dataset, a novel graph-based knowledge resource designed to enhance the logical reasoning capabilities of language models. Built using a Self-Prompting approach, the dataset combines 35 predefined relations with 133 dynamically generated relations, overcoming the limitations of fixed relational schemas. This integration results in a dual-triple structure(Head–Predefined Relation–Tail) and (Tail–Dynamic Relation–Additional Tail)that captures complex, multi-step inferences essential for advanced reasoning.

Experimental results show that models fine-tuned on the hybrid-relation Graph Dataset consistently outperform both baseline models and those fine-tuned on a general-purpose control dataset (CC News), with notable gains on Commonsense QA, Riddle Sense, and ARC-Challenge. The introduction of diverse dynamic relations, such as Analogous, Contextual, and Complementary, equips models with the flexibility to perform nuanced, context-sensitive reasoning. Performance improvements in causal and commonsense reasoning tasks further validate the dataset's effectiveness in strengthening inference skills.

With its scalability and adaptability, the hybrid-relation Graph Dataset offers a robust foundation for a wide range of reasoning-oriented applications. By advancing reasoning-focused dataset construction and refining automatic relation generation, this work contributes to narrowing the gap between general language understanding and sophisticated multi-step logical inference, paving the way for future models capable of more robust and context-aware reasoning.

REFERENCES

[1] T. B. Brown et al., "Language models are few-shot learners," 2020, *arXiv:2005.14165*.

[2] E. M. Bender, T. Gebru, A. Mcmillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?" in *Proc. ACM Conf. Fairness, Accountability, Transparency*, Mar. 2021, pp. 610–623.

[3] M. Sap, R. L. Bras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi, "ATOMIC: An atlas of machine commonsense for *if-then* reasoning," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 3027–3035.

[4] J. D. Hwang, C. Bhagavatula, R. L. Bras, J. Da, K. Sakaguchi, A. Bosselut, and Y. Choi, "(Comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, 2021, pp. 6384–6392.
[5] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Celikyilmaz, and Y. Choi, "COMET: Commonsense transformers for automatic knowledge graph construction," 2019, *arXiv:1906.05317*.
[6] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord, "Think you have solved question answering? Try ARC, the AI2 reasoning challenge," 2018, *arXiv:1803.05457*.
[7] A. Talmor, J. Herzig, N. Lourie, and J. Berant, "CommonsenseQA: A question answering challenge targeting commonsense knowledge," 2018, *arXiv:1811.00937*.
[8] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi, "HellaSwag: Can a machine really finish your sentence?" 2019, *arXiv:1905.07830*.
[9] T. Khot, P. E. Clark, M. Guerquin, P. Jansen, and A. Sabharwal, "QASC: A dataset for question answering via sentence composition," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, 2020, pp. 8082–8090.
[10] B. Yuchen Lin, Z. Wu, Y. Yang, D.-H. Lee, and X. Ren, "RiddleSense: Reasoning about riddle questions featuring linguistic creativity and commonsense knowledge," 2021, *arXiv:2101.00376*.
[11] R. E. Speer, J. Chin, and C. Havasi, "ConceptNet 5.5: An open multilingual graph of general knowledge," in *Proc. AAAI Conf. Artif. Intell.*, vol. 31, 2017, pp. 4444–4451.
[12] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor, "Freebase: A collaboratively created graph database for structuring human knowledge," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, Jun. 2008, pp. 1247–1250.
[13] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, and C. Bizer, "DBpedia—A large-scale, multilingual knowledge base extracted from Wikipedia," *Semantic Web*, vol. 6, no. 2, pp. 167–195, 2015.
[14] F. M. Suchanek, G. Kasneci, and G. Weikum, "Yago: A core of semantic knowledge," in *Proc. 16th Int. Conf. World Wide Web*, May 2007, pp. 697–706.
[15] J. Li, J. Wang, Z. Zhang, and H. Zhao, "Self-prompting large language models for zero-shot open-domain QA," 2022, *arXiv:2212.08635*.
[16] J. Achiam et al., "GPT-4 technical report," 2023, *arXiv:2303.08774*.
[17] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, 2018, p. 2.
[18] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A robustly optimized BERT pretraining approach," 2019, *arXiv:1907.11692*.
[19] P. He, X. Liu, J. Gao, and W. Chen, "DeBERTa: Decoding-enhanced BERT with disentangled attention," 2020, *arXiv:2006.03654*.
[20] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter," 2019, *arXiv:1910.01108*.



JONGWON RYU (Student Member, IEEE) received the M.S. degree in artificial intelligence from Chung-Ang University, Seoul, Republic of Korea, in 2025, where he is currently pursuing the Ph.D. degree. His research interests include natural language processing and various multimodal reasoning.



MINGI KIM (Student Member, IEEE) received the B.A. degree in business administration from Chung-Ang University, Seoul, Republic of Korea, in 2023, where he is currently pursuing the M.S. degree. His research interests include fairness, natural-language processing, and various multimodal learning.



JUNYEONG KIM (Member, IEEE) received the B.S., M.S., and Ph.D. degrees from KAIST, Republic of Korea, in 2015, 2017, and 2021, respectively. He is currently a Postdoctoral Researcher with the Artificial Intelligence and Machine Learning Laboratory, School of Electrical Engineering, KAIST. He is also an Assistant Professor with the Department of AI, Chung-Ang University, Republic of Korea. His research interests include visual–language reasoning, answering visual questions, and video-based problems.