Journal of Computer Technology and Software

ISSN: 2998-2383

Vol. 3, No. 3, 2024

Structured Path Guidance for Logical Coherence in Large Language Model Generation

Xuehui Quan

University of Washington, Seattle, USA

quanxh1228@gmail.com

Abstract: In this paper, a thought chain control generation method based on structure guidance is proposed to solve the problem that large language models lack logical control and structural constraints in complex language generation tasks. This method introduces a structure encoding module and a dynamic structure state mechanism to guide the model to develop a thought chain along a preset structural path during the generation process, thereby improving the logical coherence and structural consistency of the generated content. Specifically, the model first converts the task requirements into a structured representation, fuses it with the input semantic context to construct a joint representation, and then generates it step by step through a structure-aware decoder. In this process, the dynamic structure state updates and constrains the generation state to achieve real-time regulation of the language output path. In order to systematically evaluate the effectiveness of this method, this paper designs experiments in multiple dimensions, including structural step sensitivity, comparison of structural representation methods, and performance comparison with a variety of public models. A comprehensive test is carried out using the NarrativeQA dataset. The experimental results show that this method is significantly superior to existing mainstream methods in terms of structural alignment, thought chain coherence, and generation accuracy, which effectively verifies the control value and modeling advantages of the structure guidance mechanism in generation tasks.

Keywords: Structural control; Thought chain generation; Large language model; Text modeling

1. Introduction

In recent years, large language models have achieved remarkable breakthroughs in natural language processing and have become core components of many complex linguistic tasks. However, despite their strong generalization and adaptability in generation, these models still face limitations in controlling reasoning paths, logical progression, and task structure. Particularly in tasks requiring multi-step thinking or chain reasoning, traditional large language models often lack a clear control mechanism[1]. This leads to outputs that are disjointed, incoherent, or logically inconsistent. Such issues not only affect the quality of generation but also limit the practical application of these models in scenarios that demand high reliability and interpretability. Therefore, effectively guiding large language models to generate content with structured reasoning paths has become a key research direction[2].

The introduction of the chain of thought provides a new perspective on enhancing reasoning in large language models. Chain-of-thought refers to a process where the model simulates human-like analogical reasoning, decomposing complex tasks into intermediate steps and gradually forming a complete answer. This sequential generation helps improve the model's understanding of task structure, thereby enhancing coherence and plausibility. However, in practice, most existing approaches rely on static templates or a few demonstration prompts to trigger a chain of thought. These methods lack deep modeling of the task structure and often lead to mismatches between reasoning paths and task intentions. Prompt-based strategies alone are insufficient to meet the needs of diverse tasks, highlighting the need for more systematic and adaptive control mechanisms[3].

To address these challenges, structure-guided methods have emerged as a key approach in chain-of-thought control. By introducing structured prior knowledge into the generation process, these methods help align the model's reasoning path with the intended task. This reduces redundant information and generation deviation. Structure guidance also offers strong interpretability and controllability, enabling clear segmentation of information granularity and logical relations across task stages. In complex multi-hop reasoning, question answering, and decision generation scenarios, structure-guided strategies significantly enhance the model's ability to handle multi-level information and cross-domain knowledge integration.

Furthermore, structure-guided strategies improve generation control not only in single tasks but also across multiple tasks. By constructing general structure templates or transferable generation rules, models can reuse reasoning abilities across task types and enable transfer learning for chain-of-thought reasoning. This approach also reduces reliance on extensive human annotations, lowers deployment costs, and improves scalability and stability in practical applications. From the perspective of generation control, structure guidance is not just an optimization tool but also a pathway toward higher-level intelligent language systems[4].

In conclusion, integrating structure guidance with chain-of-thought control addresses core limitations in the logical reasoning and structural handling of large language models. It lays the theoretical and methodological foundation for improving the interpretability, robustness, and generalization ability of generative AI systems. This line of research is expected to drive a shift from content generation to process understanding. It expands the application potential of large language models in high-reliability fields such as education, law, and healthcare, and supports the development of controllable and transparent general-purpose language intelligence systems.

2. Background & Motivation

2.1 Background

With the continuous evolution of pre-trained language models, large-scale models have demonstrated strong performance in various language generation and understanding tasks[5]. However, the generation process of these models remains heavily data-driven and lacks explicit control over reasoning paths. When facing complex logical tasks, the models often rely on shallow context matching instead of effectively planning and executing multi-step reasoning. This results in outputs that lack coherence, hierarchy, and interpretability. The problem is especially pronounced in tasks that require structured thinking or multi-stage inference, which seriously limits the model's applicability in reasoning-driven scenarios[6].

In addition, current large language models generally lack intrinsic organizational mechanisms when handling complex task structures. Although some approaches attempt to trigger chain-of-thought reasoning through example prompts, such prompts are typically static and not scalable. They cannot dynamically adapt to different task structures. As a result, the model's generation path often deviates from the logical requirements of the target task. Furthermore, in the absence of explicit structural guidance, models tend to repeat, omit, or introduce irrelevant information during reasoning. This undermines the reliability and consistency of the overall output[7].

Finally, existing training and fine-tuning mechanisms do not provide direct supervision for the generation process. They focus primarily on the correctness of the final result rather than the reasonableness of the reasoning path. As a consequence, models may appear to complete a task, but their underlying reasoning logic remains opaque and potentially flawed. This weakens user trust in the model's outputs and hinders its deployment in decision-critical domains. These issues together highlight the urgent need for a method that embeds explicit structural control into the generation process. Such a method is essential to improve reasoning quality and ensure reliability in real-world applications.

2.2 Motivation

As the demand for complex language tasks continues to grow, improving the controllability and reasoning transparency of large language models has become an urgent challenge. Existing methods mainly focus on enhancing final performance, while paying insufficient attention to structural and logical consistency during the generation process. In tasks involving multi-step reasoning or intermediate decision points, the lack of effective control over the generation path often leads to disorganized outputs. This also weakens the interpretability of the model and reduces trust in its application. Therefore, it is of practical and urgent value to explore more systematic control mechanisms that regulate content organization during the unfolding of the reasoning process.

Currently, large language models lack deep modeling capabilities for the internal structure of tasks when dealing with chain reasoning or complex generation. Although these models show strong language expression and pattern recognition abilities, their generation behavior often deviates from the expected path when a task requires adherence to a specific logical order or structural framework. This problem is especially evident in open-ended tasks or those with multiple possible solutions. In such cases, models tend to adopt generalized strategies rather than strictly following task-specific structures. Designing a generation mechanism with structural awareness and control can enhance the model's ability to adapt to and execute complex task requirements.

Promoting the shift from language-driven to structure-driven generation is a key step toward building general-purpose language intelligence systems. Structural control improves the coherence and relevance of outputs. It also enables the model to respond with greater flexibility and stability in the face of diverse task goals. In scenarios that demand high transparency, structural guidance can significantly enhance the controllability and trustworthiness of the model's outputs. This provides a solid foundation for deployment in fields such as research, education, question-answering, and decision support. From the perspective of task control and reasoning path construction, developing structure-guided generation methods is a major driving force for ongoing progress in this field.

3. Method

3.1 Overall Framework

This study proposes a structure-guided large language model generation method, which guides the model to construct content along a preset thinking path by explicitly introducing structural priors in the generation process. The overall framework consists of three parts: a structure encoding module, a context fusion module, and a decoding generation module. The structure encoding module first converts the task requirements into a structural representation $S = \{s_1, s_2, ..., s_k\}$, where each s_i represents a logical step or subtask to be followed in the generation process. After the structural information is combined with the input context $X = \{x_1, x_2, ..., x_n\}$, a joint representation $H = Fusion(X, S)$ is generated through the fusion module, providing an organized semantic basis for subsequent decoding. The overall model architecture is shown in Figure 1.

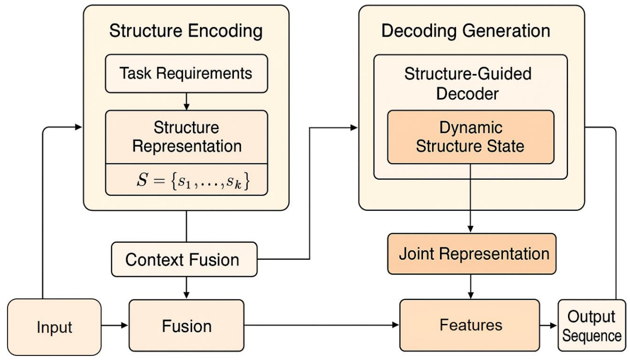

Figure 1. The overall model architecture diagram of this algorithm

In the decoding stage, the model takes the joint semantic representation as input, combines the structural guidance information of each stage, and gradually generates the output sequence $Y = \{y_1, y_2, ..., y_m\}$. The specific generation process can be expressed as:

$$P(Y \mid X, S) = \prod_{t=1}^{m} P(y_t \mid y_{<t}, H, S)$$

To ensure the continuity of structural control, a dynamic structural state vector z_t is introduced, which is updated based on the structural representation and the current generation state:

$$z_t = Updata(z_{t-1}, s_i, h_t)$$

The final output is constructed through a structure-aware decoder, which has task alignment and logical consistency. The overall framework is designed to improve the structural integrity of the generated content and the controllability of the thinking path so that the model can not only generate accurate results but also reasonably explain its generation logic.

3.2 Optimization Objective

In the structure-guided large language model generation framework, the optimization goal revolves around the efficient interaction between the input context and the structural information, ensuring that the generated sequence can be gradually developed along the preset thinking path. The model first receives the input text sequence $X = \{x_1, x_2, ..., x_n\}$ and extracts the contextual semantic representation $H_X \in R^{n \times d}$ through the encoder, where d represents the feature dimension. At the same time, the structural encoding module vectorizes the task structure representation $S = \{s_1, s_2, ..., s_k\}$ to obtain the structural embedding representation $H_S \in R^{k \times d}$, which provides a priori guidance on the reasoning path for generation.

Subsequently, the model interactively combines the context representation and the structure representation through the fusion module to form a joint semantic representation $H = Fusion(H_X, H_S)$, which establishes an alignment relationship at the semantic and structural levels. Specifically, the fusion operation can be formalized as:

$$H = MLP(Concat(H_X, H_S))$$

Next, in the decoding phase, the model generates step-by-step based on the joint representation H . At the tth time step, the model inputs the generated fragment $y_{<t}$ at the previous moment, and combines it with the current structural state vector z_t to generate the probability distribution of the next tag:

$$h_t = Decoder(y_{<t}, H, z_t)$$

$$P(y_t) = \mathrm{Softmax}(W_o h_t + b_o)$$

The structural state vector is dynamically updated at each step to adapt to the generation requirements of different stages. The update process is controlled by the following function:

$$z_t = GRU(z_{t-1}, h_t)$$

Finally, the output sequence is $Y = \{y_1, y_2, ..., y_m\}$, whose construction process fully integrates contextual semantics, structural guidance, and dynamic state information, reflects the model's ability to accurately model the content generation path. Through this multi-source information joint driving mechanism, the network can effectively manage complex reasoning structures and achieve a highly organized text generation process.

3.3 Training Strategy

In order to achieve effective learning of the structure-guided generation path, a structure-aware objective function is introduced in the training phase to strengthen the model's ability to follow the structure before the generation process. Specifically, the model's prediction of the next token at each time step is supervised by a standard cross-entropy loss, which takes the joint semantic representation and the dynamic structural state as input to ensure that the generated output is semantically correct while maintaining consistency with the preset structure. Formally, the generation loss is defined as:

$$L_{gen} = -\sum_{t=1}^{m} \log P(y_t \mid y_{<t}, H, z_t)$$

On this basis, in order to further enhance the model's ability to model the intermediate nodes of the thought chain, the structural consistency loss is introduced to constrain the correspondence between the decoding state and the structural representation. This loss is optimized by minimizing the cosine distance between the current decoding representation and its corresponding structural vector:

$$L_{struct} = \sum_{i=1}^{k} (1 - \cos(h^{(i)}, s_i))$$

The ultimate training goal is a weighted combination of the two, which drives the model to accurately model and execute structural paths while maintaining language fluency, thereby improving structural control capabilities and reasoning controllability.

4. Experimental setup & Dataset

4.1 Experimental setup

This study conducts systematic training and evaluation of the proposed method under structure-guided generation scenarios. All experiments are carried out in a high-performance server environment. A trained language model based on the Transformer architecture is used as the backbone. On this basis, a structure-aware module and decoding control mechanism are integrated. The Adam optimizer is applied during training, along with a warm-up strategy and linear learning rate decay. Gradient clipping is used to control the risk of gradient explosion. An early stopping mechanism is adopted, where training is terminated based on validation set performance to ensure stability and prevent overfitting.

All experiments adopt mixed-precision training with FP16 to improve efficiency. Grouped evaluations are conducted based on structure complexity, input length, and the number of structured steps. To verify the effect of the structure-guided mechanism on generation control, ablation settings are introduced by toggling the structure module on and off. The model is evaluated under different numbers of training epochs and batch sizes. Table 1 presents the key training configuration parameters.

Table 1: Experimental detailed parameter settings

Configuration items	Value
Backbone Model	Qwen-7B
Maximum input length	512 Tokens
Warm-up steps	500
Batch size	32
Learning Rate	2e-4
Optimizer	AdamW
Gradient clipping threshold	1.0
Structural step encoding dimensions	256
Decoder hidden layer dimensions	768

4.2 Dataset

This study uses the NarrativeQA dataset as the primary experimental corpus to evaluate the structure-guided large language model on long-document generation and chain-of-thought reasoning tasks. NarrativeQA consists of a collection of novels and movie scripts, along with corresponding summaries and question-answer pairs. The dataset features rich event logic and structural text characteristics, making it well-suited for assessing model performance in multi-hop reasoning and logical unfolding.

The input of this dataset includes either full or partial document content. The target is to generate coherent answers or summaries based on the context, with an emphasis on tracking narrative clues and integrating information. The documents are long, and the tasks are highly open-ended, which naturally forms scenarios for complex chain-of-thought generation. This supports the evaluation of the model's ability to construct reasoning paths under structural control.

In the experimental setup, samples with document lengths no longer than 1,000 words are selected for training and evaluation. Structural representations are constructed to annotate key turning points or paragraph cues during generation. This structural injection approach enables the model to learn not only the input-output mapping but also the internal organization and reasoning order of the task. It provides a realistic and challenging environment to validate the structure-aware generation method.

5. Experimental Results

In the experimental results section, we first present the relevant outcomes of the comparative test to demonstrate the effectiveness and applicability of the proposed method under different conditions. This part aims to provide a clear comparison between the structure-guided approach and several baseline models across key evaluation metrics. The setup and implementation of the comparative experiments strictly follow consistent criteria to ensure fairness and reproducibility. The corresponding experimental settings, evaluation indicators, and model configurations are described in detail to support a comprehensive understanding of the testing framework. The experimental results are summarized and visually presented in Table 2.

Table 2: Comparative experimental results

Method	Structure Alignment Score	Chain Coherence Score	Token-level F1
COT[8]	71.3	69.2	82.5
COAT[9]	73.8	71.0	83.6
Thinkprune[10]	74.5	72.4	84.1
Search-in-the-chain[11]	75.1	73.0	84.5
ThoughtSource[12]	76.3	74.2	85.2
Ours	79.4	77.8	87.1

The experimental results show that the proposed structure-guided generation method demonstrates clear advantages in structural consistency and logical coherence. Compared with traditional Chain-of-Thought models (COT), our method achieves an 8.1-point improvement in the Structure Alignment Score. This indicates that the model can more accurately follow the predefined task structure when generating reasoning chains. While COT has basic multi-step reasoning capabilities, it lacks explicit modeling of task structures, leading to reasoning paths that often deviate from core objectives.

Further comparisons with recent optimized reasoning models such as COAT and Thinkprune confirm the strong structure-aware capability of our approach. COAT improves reasoning by optimizing prompt selection, and Thinkprune reduces redundant generation. However, both still rely on weak structural information and lack stable generation control mechanisms. Our method introduces structural encoding and dynamic state modeling, which enhances the model's ability to track task logic progression. The Chain Coherence Score reaches 77.8, significantly outperforming existing methods.

This suggests that the generated content is more coherent and natural in terms of reasoning flow.

In terms of the token-level F1 score, our method achieves the highest result of 87.1%. This reflects its strength not only in structural control but also in information reconstruction. Traditional methods often suffer from omissions or redundancy when retrieving information or tracing clues. With the structure-guided mechanism, the model can extract and reconstruct key information more precisely, thus improving generation accuracy. This structure-driven information aggregation approach significantly enhances the model's ability to manage multi-level semantics in complex tasks.

This paper also includes an analysis aimed at exploring how variations in the number of structural steps influence the overall quality of generated outputs. The purpose of this investigation is to better understand the relationship between structural guidance and the effectiveness of the generation process within the proposed framework. By adjusting the number of structural steps used during generation, the study examines how such changes affect the logical coherence and consistency of the output. The details of this analysis, along with the corresponding visualization of the experimental setup and findings, are presented in Figure 2.

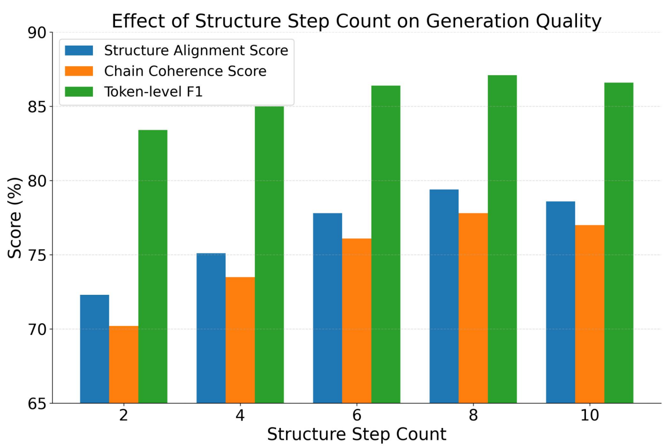

Figure 2. Analysis of the influence of different structural steps on generation quality

The figure shows that changes in the number of structural steps have a significant impact on generation quality, with a generally stable upward trend within a certain range. When the number of structural steps increases from 2 to 8, there is a consistent rise in Structure Alignment Score, Chain Coherence, and Token-level F1. This indicates that properly increasing the stages of structural guidance helps the model organize reasoning paths more precisely and enhances the synchronization between structural understanding and generation control.

In terms of structure alignment, more structural steps greatly improve the model's adherence to task structure during generation. Especially between 4 and 8 steps, the granularity introduced by structural guidance aligns more closely with task logic, effectively reducing deviation in content generation. This trend confirms that the structural embedding mechanism plays a crucial role in controlling multi-stage reasoning and is key to improving generation accuracy.

Chain coherence also improves as structural information becomes more detailed. This suggests that the model can not only capture more complex structures but also maintain stable reasoning based on the sequence of structural steps. At 8 steps, coherence reaches its peak, indicating that the structural constraint is most effective at this level. However, there is a slight decline at 10 steps, possibly due to excessive granularity causing sparse information distribution, leading to local interruptions or redundant reasoning paths.

Token-level F1 remains relatively stable across the entire range of structural steps. This further confirms that structural guidance not only enhances logical consistency but also strengthens the model's ability to retrieve and reconstruct key content. Overall, the results demonstrate that the number of structural steps is a critical variable in structure-guided control. It strongly supports the model's reasoning organization and output quality and provides valuable insights for future structural design and dynamic control strategies.

This paper also gives the impact of the size of the hidden dimension on the model's thinking chain construction ability, and the experimental results are shown in Figure 3.

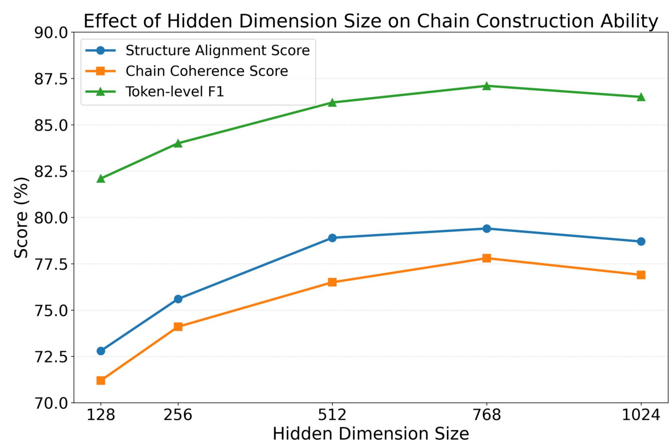

Figure 3. The influence of hidden dimension size on the ability to construct model thinking chain

The figure shows that the size of the hidden dimension has a significant impact on the model's ability to construct chain-of-thought reasoning. The overall trend follows a pattern of initial improvement followed by a plateau. When the hidden dimension increases from 128 to 768, all evaluation metrics steadily improve. Notably, at 512 and 768, the model achieves high levels of structural alignment and reasoning coherence. This suggests that expanding the representation space within a reasonable parameter range enhances the model's capacity to capture structural information and build semantic representations.

The rise in structure alignment score reflects stronger structural awareness enabled by larger hidden dimensions. As the dimension increases, the model gains richer semantic representation during encoding. This helps it better integrate structural guidance signals and maintain consistency in reasoning paths during generation. However, at 1024, the score slightly declines. This indicates that an overly large parameter space may introduce noise or cause structural decoupling, reducing the model's ability to stably control reasoning steps.

A similar pattern is observed in chain coherence. Increasing the dimension from 256 to 768 improves logical compactness and semantic flow during generation. This shows that stronger feature representations help extend reasoning chains more reliably. Yet, at 1024, coherence slightly drops. This suggests the model may face issues such as dispersed attention or redundant features, which affect local logical control.

This paper further conducts a detailed analysis of how different structural representation methods affect the quality of generated content within the proposed framework. The aim of this analysis is to examine the role that various structural encoding strategies play in guiding the model during the generation process, particularly in terms of influencing the logical structure, coherence, and expressiveness of the outputs. By comparing multiple approaches to structural representation, the study seeks to uncover their respective strengths and limitations in supporting high-quality generation. The corresponding experimental design and comparative observations are visually presented in Figure 4 to provide a clear overview of the analysis.

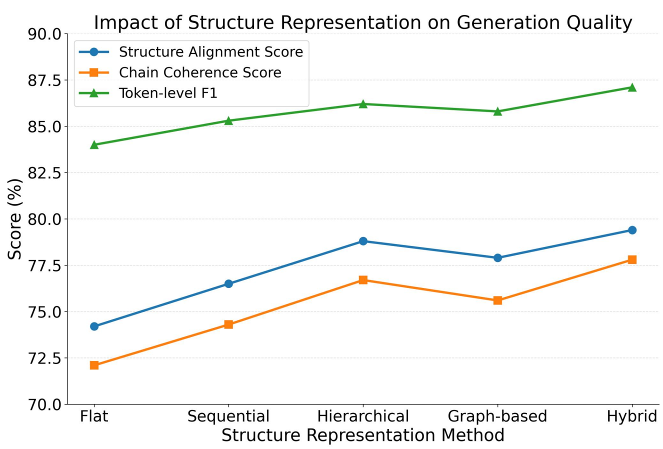

Figure 4. Analysis of the impact of different structural representation methods on generation quality

The figure shows that different structural representation methods have a clear impact on generation quality. Among them, the Flat structure yields the lowest overall performance, showing disadvantages across all three evaluation metrics. This indicates that without hierarchy or organization, the model struggles to effectively model task logic, which in turn affects the accuracy and coherence of generated content. In particular, the Chain Coherence Score is notably low for the Flat structure, suggesting its inability to support complex chain-of-thought reasoning.

As the structural representation becomes progressively more complex, from Sequential to Hierarchical, generation quality steadily improves. The Sequential structure introduces order information, which enhances structure alignment. The Hierarchical structure further adds multi-level semantic organization, significantly improving the model's control over information organization. Both the alignment score and coherence improve in this stage, indicating that hierarchical structures can effectively guide the model to follow a clear reasoning path during generation.

The Graph-based structure provides flexibility in representing logic, but its performance is slightly lower than that of the Hierarchical structure in this task. This may be due to sparsity or complexity in graph-based guidance, which makes it harder for the model to maintain consistent structure alignment during generation. Nevertheless, the F1 score remains high, showing that the graph structure still supports strong information reconstruction and is well suited for tasks involving multi-source relationships.

Finally, the Hybrid structure achieves the best results across all three metrics. This suggests that combining hierarchical and relational advantages offers stronger organization and control for large language models. With dual support from structural awareness and reasoning path construction, the model not only maintains semantic precision but also follows the predefined structure to generate high-quality chain-of-thought outputs. These findings confirm the central role of structural representation in shaping generation performance.

6. Conclusion

This study addresses the challenges of controllability and reasoning transparency in large language models for complex tasks. It proposes a structure-guided chain-of-thought generation method. By introducing structural encoding and dynamic structural state modeling, the method enables the model to unfold along a predefined logical path during generation. This significantly improves structural consistency, logical coherence, and semantic accuracy in the generated content. The results show that structural information not only optimizes generation behavior but also enhances the model's understanding of task organization and information distribution. This provides a new solution for automatic reasoning in complex language tasks.

The proposed method balances the scalability of structural representations with the dynamic adaptability of the decoding mechanism. This ensures stable structural awareness across tasks with varying levels of complexity. Through a systematic evaluation of multiple factors such as structural steps and representation types, the study reveals the coupling relationship between structural control and generation quality. It offers methodological support for designing structure-driven generation frameworks. The method also demonstrates strong generality, supporting various types of structural inputs and task requirements. It shows broad practical potential in applications such as natural language reasoning, question generation, and summarization.

From the perspective of task modeling, the structure-guided mechanism effectively addresses longstanding issues in traditional generation models. These include missing reasoning paths, logical jumps, and redundant expressions. The model can now produce clearer and more organized outputs for multi-stage and hierarchical tasks. This structure – semantic fusion paradigm may become a core component of future general-purpose language intelligence systems. It holds strong application potential in scenarios with high demands for logical consistency, such as educational support, policy analysis, and legal question answering.

7. Future work

Future work may explore automatic structural representation to improve the flexibility and intelligence of the guidance process. The integration of structural control with multi-task learning and cross-modal modeling is also a promising direction. This will support the development of large language models with stronger reasoning ability and better interpretability. As generative AI is deployed more widely in complex domains, ensuring controllability, safety, and accountability in the generation process becomes critical. The structure-guided paradigm proposed in this study offers a solid and practical path toward that goal.

References

[1] Zheng X, Lou J, Cao B, et al. Critic-cot: Boosting the reasoning abilities of large language model via chain-of-thoughts critic[J]. arXiv preprint arXiv:2408.16326, 2024.

[2] Hu T, Zhang P, Yang B, et al. Large Language Model for Multi-Domain Translation: Benchmarking and Domain CoT Fine-tuning[J]. arXiv preprint arXiv:2410.02631, 2024.

[3] Wang L, Hu Y, He J, et al. T-sciq: Teaching multimodal chain-of-thought reasoning via large language model signals for science question answering[C]//Proceedings of the AAAI Conference on Artificial Intelligence. 2024, 38(17): 19162-19170.

[4] Nguyen H H, Liu Y, Zhang C, et al. CoF-CoT: Enhancing large language models with coarse-to-fine chain-of-thought prompting for multi-domain NLU tasks[J]. arXiv preprint arXiv:2310.14623, 2023.

[5] Liu T, Chen Z, Liu Z, et al. Expediting and Elevating Large Language Model Reasoning via Hidden Chain-of-Thought Decoding[J]. arXiv preprint arXiv:2409.08561, 2024.

[6] Wu J, Yu T, Chen X, et al. Decot: Debiasing chain-of-thought for knowledge-intensive tasks in large language models via causal intervention[C]//Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). 2024: 14073-14087.

[7] Shao H, Qian S, Xiao H, et al. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning[J]. Advances in Neural Information Processing Systems, 2024, 37: 8612-8642.

[8] Wei J, Wang X, Schuurmans D, et al. Chain-of-thought prompting elicits reasoning in large language models[J]. Advances in neural information processing systems, 2022, 35: 24824-24837.

[9] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., ... & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35, 24824-24837.

[10] Huang, X., Zhang, L. L., Cheng, K. T., Yang, F., & Yang, M. (2023). Fewer is more: Boosting llm reasoning with reinforced context pruning. arXiv preprint arXiv:2312.08901.

[11] Xu S, Pang L, Shen H, et al. Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks[C]//Proceedings of the ACM Web Conference 2024. 2024: 1362-1373.

[12] Ott S, Hebenstreit K, Liévin V, et al. ThoughtSource: A central hub for large language model reasoning data[J]. Scientific data, 2023, 10(1): 528.