Beyond Chain-of-Thought: A Survey of Chain-of-X Paradigms for LLMs

Yu Xia[1,2] **Rui Wang**[3] **Xu Liu**[1] **Mingyan Li**[1] **Tong Yu**[4]
Xiang Chen[4] **Julian McAuley**[2] **Shuai Li**[1*]
[1]Shanghai Jiao Tong University [2]UC San Diego [3]Duke University [4]Adobe Research
{yux078, jmcauley}@ucsd.edu {tyu, xiangche}@adobe.com
rui.wang16@duke.edu {liu_skywalker, QYLJM1217, shuaili8}@sjtu.edu.cn

Abstract

Chain-of-Thought (CoT) has been a widely adopted prompting method, eliciting impressive reasoning abilities of Large Language Models (LLMs). Inspired by the sequential thought structure of CoT, a number of Chain-of-X (CoX) methods have been developed to address challenges across diverse domains and tasks. In this paper, we provide a comprehensive survey of Chain-of-X methods for LLMs in different contexts. Specifically, we categorize them by taxonomies of nodes, i.e., the X in CoX, and application tasks. We also discuss the findings and implications of existing CoX methods, as well as potential future directions. Our survey aims to serve as a detailed and up-to-date resource for researchers seeking to apply the idea of CoT to broader scenarios.

1 Introduction

Large Language Models (LLM) have shown strong reasoning capabilities when prompted with the Chain-of-Thought (CoT) method (Wei et al., 2022). The essence of CoT is to decompose complex problems into sequences of intermediate subtasks (Chu et al., 2024). By handling these subtasks step by step, LLMs are able to focus on important details and assumptions, which substantially improves the performance across a wide range of reasoning tasks (Huang and Chang, 2023; Chu et al., 2024). CoT's intermediate steps can also offer a more transparent reasoning process, facilitating easier interpretation and evaluation of LLMs (Yu et al., 2023b).

With the success of CoT, a number of Chain-of-X (CoX) methods have subsequently been developed (Yu et al., 2023a). Extending beyond reasoning thoughts, recent CoX methods have constructed the chain with various components, such as Chain-of-Feedback (Lei et al., 2023; Dhuliawala et al., 2023), Chain-of-Instructions (Zhang et al., 2023d; Hayati et al., 2024), Chain-of-Histories (Luo et al., 2024; Xia et al., 2024e), etc. These methods have been applied to diverse tasks involving LLMs beyond reasoning, including multi-modal interaction (Xi et al., 2023a; Zhang et al., 2024a), hallucination reduction (Lei et al., 2023; Dhuliawala et al., 2023), planning with LLM-based agents (Zhang and Zhang, 2024; Zhang et al., 2024b), etc.

Motivation Despite their growing prevalence, these CoX methods have not yet been collectively examined or categorized, leaving a gap in our understanding of their potential. This survey offers a structured overview capturing CoX's essence and diversity for further exploration and innovation.

Distinguishing Focus While several surveys have explored CoT (Chu et al., 2024; Yu et al., 2023b; Besta et al., 2024), they focus primarily on the reasoning thoughts of different structures, e.g., Chain-of-Thought as illustrated in Figure 1(a). In contrast, this paper focuses on the multifaceted component designs of Chain-of-X beyond reasoning thoughts as shown in Figure 1, offering insights of the CoT concept in broader domains. We present a comprehensive review by taxonomies of the X in CoX and tasks to which these methods are applied.

Overview of the Survey We first provide background information on Chain-of-Thought and define Chain-of-X as its generalization (§2). Next, we categorize CoX methods by the types of components used to construct the chains (§3). Furthermore, based on the application areas of these CoX methods, we categorize them by tasks (§4). Then, we discuss insights from existing CoX methods and explore potential future directions (§5). A detailed structure of the survey is presented in Figure 2.

2 What is Chain-of-X?

In this section, we introduce some background information about Chain-of-Thought prompting and then define a generalized concept of Chain-of-X.

* Corresponding author.

Proceedings of the 31st International Conference on Computational Linguistics, pages 10795–10809
January 19–24, 2025. ©2025 Association for Computational Linguistics

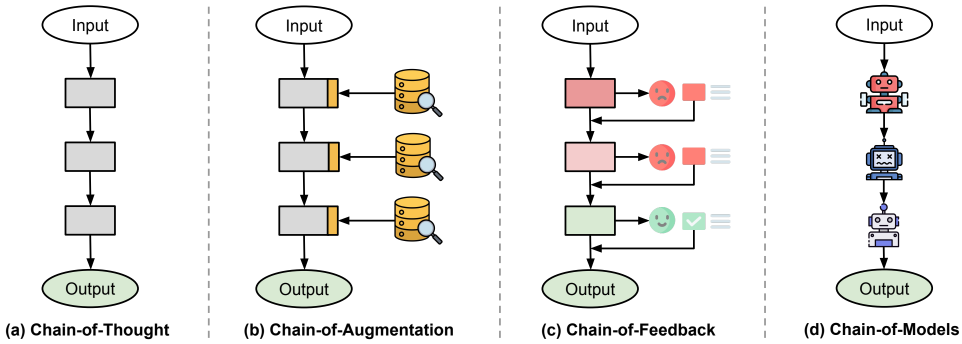

Figure 1: Illustrations of Chain-of-X paradigms with four types of nodes: (a) Intermediates, e.g., Thought (§3.1), (b) Augmentation (§3.2), (c) Feedback (§3.3), and (d) Models (§3.4).

Chain-of-Thought CoT prompting is a methodology that substantially enhances the reasoning capabilities of LLMs. Introduced by Wei et al. (2022), CoT involves prompting LLMs with a structured format of `<input, thoughts, output>`, where 'thoughts' encompass coherent and intermediate natural language reasoning steps leading to the final answer. CoT's effectiveness is most pronounced in tasks that require complex reasoning. Traditional few-shot learning methods often falter in such scenarios, as they tend to provide direct answers without rationales. In contrast, CoT prompting excels by breaking down a complex task into manageable intermediate steps. These steps guide the model through a logical progression, enhancing its capability to tackle complex problems such as arithmetic, commonsense, and symbolic reasoning (Wang et al., 2023f; Lyu et al., 2023). Additionally, Kojima et al. (2022) have also demonstrated strong performance of zero-shot CoT by prompting "*Let's think step by step.*". The explicit reasoning steps also provide a transparent pathway for the model's thought process, allowing for further evaluations and corrections (Yu et al., 2023b).

Chain-of-X Inspired by the nature of the sequential breakdown, a substantial number of CoX methods have been developed recently (Yu et al., 2023a). Here, we define CoX as a generalization of the CoT method for diverse tasks beyond LLM reasoning. We refer to the X in CoX as the 'node' of the chain structure. Beyond the thoughts in CoT prompts, the X in CoX can take various forms tailored to specific tasks, including intermediates (§3.1), augmentation (§3.2), feedback (§3.3), and even models (§3.4), as illustrated in Figure 1. We summarize the types of nodes in existing CoX methods in Figure 2. The idea of CoX is to construct a sequence of problem-related components that either compositionally contribute to the solution or iteratively refine the outputs. We define a similar structured format for CoX as `<input,` $\mathtt{X}_1$`, ...,` $\mathtt{X}_n$`, output>` where n is the length of the chain. Note that this format extends beyond prompting-based strategies like CoT and can be adapted to a variety of algorithmic frameworks or structures for diverse tasks involving LLMs. For instance, Chain-of-Verification (Dhuliawala et al., 2023) is a hallucination reduction framework that employs an LLM to generate initial responses, composes a sequence of verification questions, and revises its previous responses based on these questions. In addition to hallucination reduction, CoX methods have been applied to a variety of tasks, as shown in Figure 2, including multi-modal interaction (§4.1), factuality & safety (§4.2), multi-step reasoning (§4.3), instruction following (§4.4), LLMs as Agents (§4.5), and evaluation tools (§4.6).

3 Chain-of-X Nodes

In this section, we survey existing CoX methods by taxonomy of nodes, categorizing them as shown in Figure 2 based on the distinct nature of the nodes.

3.1 Chain-of-Intermediates

Building on the concept of utilizing intermediate steps, a natural evolution of CoT involves generalizing reasoning thoughts to other types of intermediate components. Based on the primary focuses, we further divide them into the following subtypes.

Problem Decomposition In problem decomposition, the intermediate steps consist of manageable subtasks derived from an original complex problem, which is exemplified by the classic Chain-of-Thought prompting (Wei et al., 2022) for reasoning tasks. To overcome the challenge of easy-to-hard generalization, Zhou et al. (2023) further

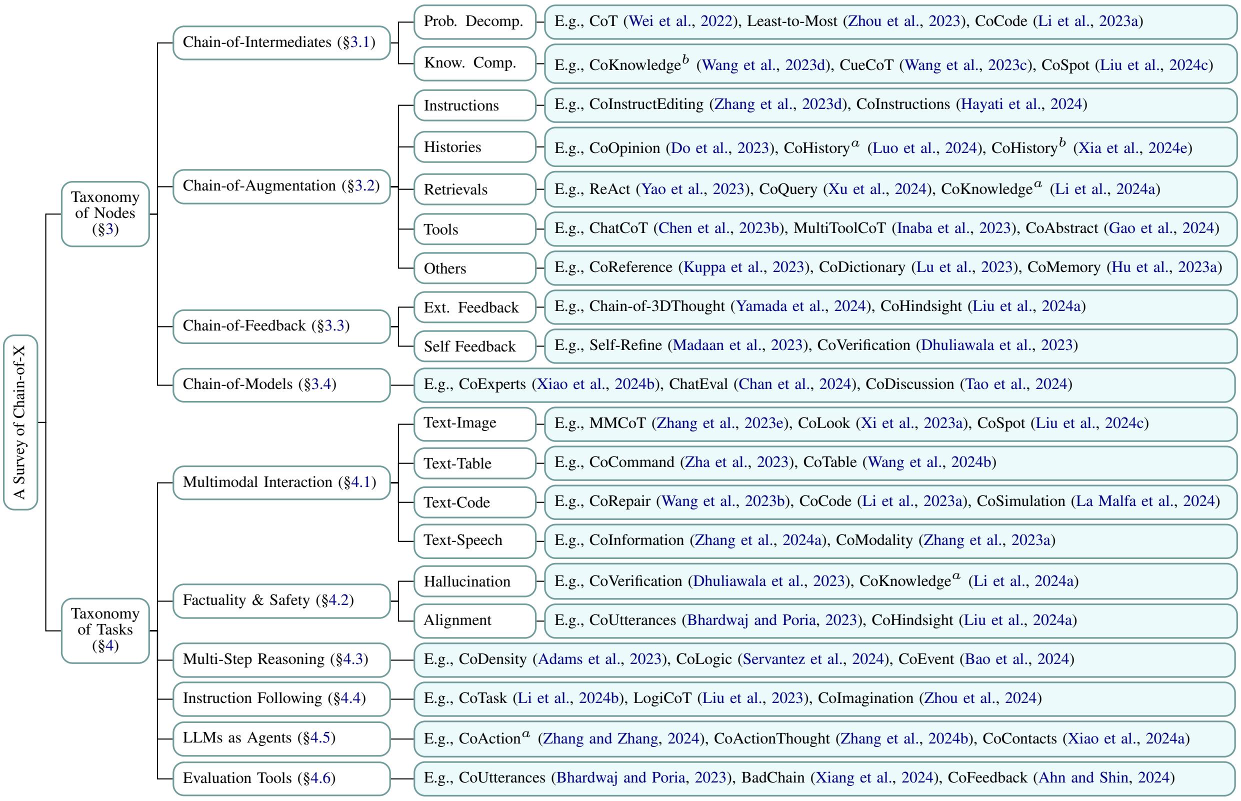

Figure 2: A survey of Chain-of-X by taxonomies of nodes and tasks (only representative methods are listed due to space limitation and a more complete version can be found in Appendix A).

introduce Least-to-Most prompting which breaks down a complex problem into simpler subtasks and solves them in sequence. Extending beyond natural languages, Chain-of-Code (Li et al., 2023a) takes advantage of the syntactic structure and precise computation of code by segmenting a complex task into programmatic subtasks, enhancing the reasoning process through simulated code outputs. Similarly motivated by computation accuracy, Chain-of-Logic (Servantez et al., 2024) applies a logical decomposition transforming rule-based reasoning tasks into a series of simple logical expressions.

While such decomposition is widely applied in reasoning tasks, the concept is also echoed in other tasks. For example, Chain-of-Event (Han et al., 2024) simplifies multi-document summarization into discrete event extraction tasks, significantly enhancing the summarization quality and factuality. Chain-of-Table (Wang et al., 2024b) restructures complex tables into question-specific formats via a sequence of strategic operations, making the data more accessible and tailored to the inquiry.

Knowledge Composition In knowledge composition, the primary goal of the intermediate steps is not simplification but the accumulation of relevant information. This approach aims to enrich the solution with a depth of understanding and details. For instance, to handle unfactual rationales generated with CoT prompting, Wang et al. (2023d) propose Chain-of-Knowledge^b to elicit LLMs to generate explicit knowledge evidence at each reasoning step for more grounded question-answering. Similarly in dialogue systems, CueCoT (Wang et al., 2023c) collects linguistic cues with intermediate steps to capture contextual user status for more personalized and engaging conversation.

Besides natural language tasks, this technique is also useful in knowledge-intensive visual tasks that require capturing specific visual details. For example, Chain-of-Spot (Liu et al., 2024c) and Chain-of-Reasoning (Uehara et al., 2024) enable vision-language models to focus on key regions of interest, improving reasoning performance with detailed visual evidences. Likewise, CCoT (Mitra et al., 2024) utilizes scene-graph representations to extract compositional knowledge from a large multimodal model, which is further used to facilitate its own response generation on vision-language tasks.

3.2 Chain-of-Augmentation

While Chain-of-Intermediates method has proven effective, it falls short when LLMs have limited knowledge for specific tasks or domains. As a result, Chain-of-Augmentation has become a popular variant of CoX, where the chain is augmented with additional knowledge. Based on the types of augmented data, we categorize them as follows.

Instructions Given a complex task, determining the next step can be nontrivial for LLMs even with few-shot CoT exemplars, due to misinterpretation or ambiguity (Zha et al., 2023). Instructions then serve as an important augmentation, guiding LLMs through complex reasoning steps or task execution processes. For instance, Chain-of-InstructEditing (Zhang et al., 2023d) harnesses this concept by generating sequential instructions to guide image editing tasks, illustrating how specific operations can refine the output for more precise editing. To avoid ambiguous user queries in table manipulation, Zha et al. (2023) introduce Chain-of-Command framework. Inferring from user instructions, it enables LLMs to employ a series of pre-defined commands for more precise table execution. In the realm of e-commerce, Li et al. (2024b) implement a similar approach Chain-of-Task, which breaks down customer interactions into manageable atomic tasks with domain-specific e-commerce instructions, significantly simplifying complex user queries.

Recently, Hayati et al. (2024) propose Chain-of-Instructions. Different from previous methods using pre-defined or human-crafted instructions, this framework iteratively uses outputs of previous steps as instructions for the next step. An instruction dataset generated is then used for fine-tuning LLMs to handle complex instructions composed of multiple subtasks. The results show that stepwise guidance can effectively improve the process and the outcomes of complex problem-solving tasks.

Histories Augmenting LLMs with historical data is essential for predictive modeling, which introduces another facet of Chain-of-Augmentation drawing contextual insights from the past. This approach is exemplified by Chain-of-Opinion (Do et al., 2023), which analyzes historical user opinions to predict future reactions, offering valuable foresight into user sentiments. In user-interface exploration, Chain-of-Action[a] (Zhang and Zhang, 2024) framework leverages past actions to guide future interactions, thereby optimizing user experience through predicted behaviors. Ma et al. (2023) take a similar approach in gaming environments like StarCraft II, where Chain-of-Summarization framework provides strategic gameplay recommendations based on a synthesis of past observations.

In addition to user modeling, the prediction of taxonomy structures also benefits from historical data, as seen in Chain-of-Layer (Zeng et al., 2024), which builds upon previously identified categories. Temporal knowledge graphs receive a forward-looking treatment as well with methods like Chain-of-History[a] (Luo et al., 2024) and Chain-of-History[b] (Xia et al., 2024e), where historical graph structures inform LLM predictions about future nodal linkages or interactions.

Retrievals As the knowledge learned from pretraining data is limited and often outdated, LLMs frequently need to acquire external knowledge. Therefore, retrieval has become a crucial aspect of Chain-of-Augmentation. As one-step retrieval is often insufficient for complex tasks, methods have been developed to intersperse reasoning chains with explicit retrievals, thereby enhancing the quality of answers. For example, ReAct (Yao et al., 2023) synergizes reasoning and acting by adaptively retrieving external knowledge to augment the reasoning chains. While ReAct prompts LLMs to make decisions, Verify-and-Edit (Zhao et al., 2023) decides when to retrieve based on less-than-majority agreement consistency and corrects erroneous reasoning chains to produce a more interpretable CoT. Further refining this concept, Li et al. (2024a) develop Chain-of-Knowledge[a], which dynamically pulls relevant information from both unstructured and structured knowledge sources, e.g., Wikidata and tables. While ReAct and Verify-and-Edit keep all retrieved information in the chain, Chain-of-Knowledge[a] makes progressive corrections and only incorporates verified retrieval results to avoid propagating misleading information.

Different from the adaptive retrieval in previous methods, another line of work explores how to compose informative queries during intermediate steps for knowledge retrieval. Press et al. (2023) propose Self-Ask, prompting LLMs to ask follow-up questions themselves and answer these sub-questions with a Google Search API before generating the final response. Similarly, IRCoT (Trivedi et al., 2023), Chain-of-Question (Huang et al., 2024), and RAT (Wang et al., 2024a) augment each intermediate step with retrieved external knowledge iteratively refining the generations. These methods directly insert retrievals into the reasoning chains, where LLMs can only reason about a local subquestion in each generation. Thus, when there

Chain-of-Retrievals	Self-Generated Query	Adaptive Retrieval	Retrieval Verification	Knowledge Sources
Self-Ask (Press et al., 2023)	✓			Textual Corpus
ReAct (Yao et al., 2023)	✓	✓		Textual Corpus
Verify-and-Edit (Zhao et al., 2023)	✓	✓		Textual Corpus
CoKnowledge[a] (Li et al., 2024a)		✓	✓	Textual & Tabular Data
IRCoT (Trivedi et al., 2023)	✓			Textual Corpus
CoQuery (Xu et al., 2024)	✓		✓	Textual Corpus
CoAction[b] (Pan et al., 2024)	✓		✓	Textual & Market Data
RAT (Wang et al., 2024a)	✓			Textual Corpus
ToG (Sun et al., 2024)		✓	✓	Knowledge Graphs
GraphCoT (Jin et al., 2024)	✓	✓		Knowledge Graphs

Table 1: A comparison of representative Chain-of-Retrievals methods from method and data source perspectives.

is a misleading sub-question, the entire reasoning chain afterwards will be affected. To address this limitation and ensure the coherence of the global chain, Xu et al. (2024) develop the Chain-of-Query framework, which can interactively revisit previous retrievals and make necessary reasoning direction adjustments by verifying retrieved results. Based on methodology designs and retrieval sources, we make further comparisons of previously discussed methods in Table 1.

Tools Besides deploying a retriever to access external knowledge, recent methods have also explored utilizing other domain specific tools. MultiToolCoT (Inaba et al., 2023) specifies available tools in the prompt and provides demonstration examples to enable tool using during CoT prompting. To achieve more natural tool using, ChatCoT (Chen et al., 2023b) models CoT reasoning as multi-turn conversations, enabling LLMs to freely interact with tools through chatting. To further improve the efficiency of interconnected tool calls, Gao et al. (2024) develop Chain-of-Abstraction, training LLMs to decode reasoning chains with abstract placeholders to be filled with knowledge from tools. Such abstract chains enable LLMs to perform decoding and calling of external tools in parallel, thus reducing the inference delay of tool responses.

Others Other domain-specific augmentation methods have also been explored, e.g., Chain-of-Empathy (Lee et al., 2023b) for empathetic response generation, Chain-of-Reference (Kuppa et al., 2023) for complex legal inquiries, and Chain-of-Dictionary (Lu et al., 2023) for machine translation. These methods not only broaden the operational scope of LLMs but also underscore the potential of domain-specific CoX enhancements.

3.3 Chain-of-Feedback

Chain-of-Feedback represents another variant of CoX. Unlike augmentation which typically precedes generation, feedback is interlaced throughout the generation process to enhance and fine-tune responses. Based on the feedback source, we categorize them as external and self feedback.

External Feedback External feedback provides valuable perspectives overlooked by LLMs themselves. For instance, to generate 3D objects that LLMs may not have seen before, Chain-of-3DThought (Yamada et al., 2024) utilizes external critiques to help iteratively hone an LLM's understanding of 3D spaces for unconventional objects. In addition to model feedback, human feedback is another important type of external feedback especially towards aligning LLMs with human preferences. Despite the success of RLHF (Ouyang et al., 2022), Chain-of-Hindsight (Liu et al., 2024a) further transforms direct human preference data into natural language feedback that better aligns with how LLMs process textual information. Such feedback allows for more precise refinement to model's outputs, ensuring that responses are both accurate and contextually appropriate.

Self Feedback Though external feedback is critical, its costs and potential unavailability have led to a growing interest in self-refinement approaches (Lee et al., 2023a). Highlighted by Madaan et al. (2023), Self-Refine first generates an initial response using an LLM and then uses the same LLM to provide feedback and refine its own response iteratively. Without any additional training, Self-Refine generates considerably better responses than direct generation. Echoing this approach, Dhuliawala et al. (2023) introduce Chain-of-Verification. Instead of directly asking LLMs to provide feedback on their own responses, Chain-of-Verification asks LLMs to first plan a series of verification questions based on the initial responses and then answer these questions themselves. After the self-assessment, LLMs are then asked to generate their

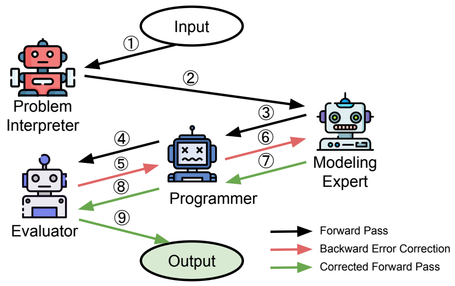

Figure 3: A simplified illustrative workflow of Chain-of-Experts (Xiao et al., 2024b).

final verified answers. Chain-of-NLI (Lei et al., 2023) adopts similar framework but formulates a series of natural language inference problems to be answered. Similar concept has also been applied in other tasks. Chain-of-Density (Adams et al., 2023) allows LLMs to iteratively refine generated summaries by incorporating self-detected missing information. Chain-of-SelfRevisions (Le et al., 2024) improves modular code generation by reusing previously generated code modules.

3.4 Chain-of-Models

Previous CoX methods have mostly been designed for a single LLM. Recognizing that different LLMs may have different specialties (Xiao et al., 2024b; Xia et al., 2024b), another line of work proposes constructing a chain of models to leverage distinct strengths of each model. Chain-of-Experts (Xiao et al., 2024b) exemplifies this collaborative strategy. As illustrated in Figure 3, it involves a consortium of expert LLMs that work in sequence, each contributing its specialized knowledge to build upon the results developed by its predecessors. This method is particularly effective in addressing intricate problems in operations research, where the complexity often exceeds the processing capabilities of a single LLM. Similarly, Qiu et al. (2024) deploy a chain of specialized LoRA (Hu et al., 2022) networks, each fine-tuned to effectively handle different domains of a broader problem. This approach ensures that specific tasks benefit from the most relevant expertise, enhancing overall efficiency and outcome accuracy. In parallel, ChatEval (Chan et al., 2024) and Chain-of-Discussion (Tao et al., 2024) employ multiple LLMs engaging in a dialogue, critiquing and refining each other's contributions before reaching a consensus in the final response. This process ensures that the synthesized output is not only comprehensive but also critically evaluated from multiple perspectives.

4 Chain-of-X Tasks

CoX can be of various forms, enabling their applications in diverse areas. This section surveys CoX methods categorized by tasks as shown in Figure 2.

4.1 Multi-Modal Interaction

Real-world problem-solving tasks often involve modalities other than text. Thus, the success of CoT has drawn attention to designing frontier CoX methods for challenges in multi-modality.

Text-Image To handle rich features from both texts and images, the knowledge composition ability of CoX methods has been crucial in capturing key information. While MultimodalCoT (Zhang et al., 2023e) incorporates image information into textual rationale generation, it lacks interaction between modalities. To address this, methods have explored using intermediate reasoning steps to infer and extract key visual information before generating final responses. For example, DDCoT (Zheng et al., 2023) utilizes structured logical chains to explicitly guide the understanding of relevant image regions. Chain-of-Look (Xi et al., 2023a) constructs a visual semantic reasoning chain based on textual cues for visual entity recognition. Chain-of-Manipulation (Qi et al., 2024) and Chain-of-Spot (Liu et al., 2024c) adopt a step-wise refinement process for identifying critical image details.

Text-Table Unlike texts, structured tabular data has been a challenging source for LLMs to reason with or manipulate (Fang et al., 2024). To this end, CoX methods show advantages in decomposing complex table operations into a sequence of manageable subtasks. Zha et al. (2023) utilize a sequence of pre-defined commands to execute table operations step by step until the queried information is found. Taking a step further, Chain-of-Table (Wang et al., 2024b) directly leverages tabular data as a part of the reasoning chain. Here, tables are not just data sources but act as evolving entities within the reasoning process, dynamically being updated in response to the LLM's queries and tasks. This iterative process allows the model to engage with the table more naturally and effectively, leading to a more nuanced understanding and manipulation of table information.

Text-Code The nature of sequential execution of code generation makes it another task benefiting from CoX methods (Zan et al., 2023). Chain-of-Repair (Wang et al., 2023b) draws inspiration from traditional debugging processes. It employs a teacher model to interpret compiler feedback and compose a chain of code repairing steps, teaching a student model to generate debugged code. Chain-of-SelfRevisions (Le et al., 2024) explores modular code generation. This method iteratively extracts and clusters sub-modules from previous generations and adds them to new reasoning chains, naturally encouraging code reuse and efficiency.

Text-Speech The field of speech generation has also seen innovative applications of CoX methods. Chain-of-Information (Zhang et al., 2024a) enhances speech synthesis by separating and then reassembling semantic and perceptual components, which allows for more nuanced and accurate speech output. Chain-of-Modality (Zhang et al., 2023a) merges textual and vocal instructions to guide speech generation. This method not only enhances the quality of speech generation but also enables LLMs to handle conversational nuances, effectively bridging the gap between textual and speech data.

4.2 Factuality & Safety

Ensuring factuality and safety in LLM outputs has been critical for practical applications (Wang et al., 2023g). Recent studies have also explored CoX for both hallucination reduction and alignment.

Hallucination Reduction LLMs have shown a propensity for generating hallucinations (Agrawal et al., 2023; Xia et al., 2024c). Two main sources of hallucinations include: i) LLMs being overconfident in their incorrect understanding of the problem and overlooking details, and ii) LLMs having limited knowledge about certain tasks and generating uncertain answers (Zhang et al., 2023c). For the first type, step-wise verification and iterative refinement have been applied to guide LLMs to reassess their initial response and focus on details, exemplified by Self-Refine (Madaan et al., 2023), Chain-of-NLI (Lei et al., 2023), and Chain-of-Verification (Dhuliawala et al., 2023). For the second type, it is essential to augment LLMs with additional knowledge grounding their responses. Several CoX methods, e.g., Chain-of-Note (Yu et al., 2023a), Chain-of-Knowledge[a] (Li et al., 2024a), and Chain-of-Action[b] (Pan et al., 2024) retrieve domain-specific knowledge at each step, effectively reducing the occurrence of unfactual generations.

Alignment Aligning LLMs with human preferences is critical to ensure that LLMs generate helpful and harmless responses. Despite the wide adoption of RLHF for LLM alignment (Ouyang et al., 2022; Xia et al., 2024d), there are still challenges including the high cost of human annotation and imperfect reward functions. To help LLMs learn from any feedback form, Liu et al. (2024a) proposes transforming preference data into a sequence of natural language sentences for supervised finetuning. Leveraging the language comprehension capabilities of LLMs, Chain-of-Hindsight achieves better alignment performance compared to RLHF. Meanwhile, Chain-of-Utterance prompting (Bhardwaj and Poria, 2023) has been proposed for LLM red-teaming. It adopts a sequential structure to establish a jailbreaking conversation between a harmful LLM and a helpful but unsafe LLM, exposing safety issues of LLMs to be aligned.

4.3 Multi-Step Reasoning

Multi-step reasoning typically demands a robust understanding of context and logic (Wei et al., 2022). These tasks require breaking down complex problems into a series of smaller, interconnected steps, building upon each step to reach a logical conclusion. The sequential nature of CoX makes it ideally suited for these tasks, including rule-based reasoning (Servantez et al., 2024), database reasoning (Hu et al., 2023a), legal reasoning (Kuppa et al., 2023), user behavior reasoning (Do et al., 2023; Han et al., 2024), graph reasoning (Zeng et al., 2024; Luo et al., 2024; Xia et al., 2024e), as well as reasoning for summarization (Adams et al., 2023; Bao et al., 2024) and machine translation (Lu et al., 2023). These varied applications demonstrate CoX's advantages in enhancing LLMs' ability to process information more effectively.

4.4 Instruction Following

Instruction following has been a celebrated ability of LLMs (Zhang et al., 2023b). The evolution of CoX methods has also led to various approaches for enhancing this feature. A notable line of works, such as Chain-of-Task (Li et al., 2024b), LogiCoT (Liu et al., 2023) and Chain-of-Imagination (Zhou et al., 2024), construct sequences of instructions for prompting or instruction tuning to handle complex tasks that require explicit step-wise guidance. While training a single LLM to follow different instructions can be costly, Chain-of-LoRA (Qiu et al., 2024) adopts a series of LoRA networks to specialize in instruction handling. After identifying an instruction type, Chain-of-LoRA applies task-specific LoRA networks to the base LLM to accomplish the respective tasks.

4.5 LLMs as Agents

The planning abilities have made LLMs strong agents across a wide range of tasks (Xi et al., 2023b). CoX methods have been explored to further boost the planning abilities of LLM-based agents. In this vein, Chain-of-Action[a] (Zhang and Zhang, 2024) and Chain-of-ActionThought (Zhang et al., 2024b) utilize a series of planned actions to guide the decision-making of agents, ensuring each step is informed by the previous actions. As discussed previously, LLM-based agents can also be augmented with historical data, e.g., Chain-of-Summarization (Ma et al., 2023), and external model feedback, e.g., Chain-of-3DThought (Yamada et al., 2024). LLMs also serve as planners in human-scene interaction tasks with Chain-of-Contacts (Xiao et al., 2024a), and in tool using with Chain-of-Abstraction (Gao et al., 2024). Chain-of-Models also naturally involves multi-agent settings such as Chain-of-Discussion (Tao et al., 2024). These methods highlight the integration of CoX in enhancing the multi-dimensional abilities of LLMs as autonomous and collaborative agents.

4.6 Evaluation Tools

Evaluating LLMs has become increasingly challenging as they grow more sophisticated (Chang et al., 2023), making CoX methods a valuable asset for evaluation purposes. Chain-of-Utterances prompting (Bhardwaj and Poria, 2023) exposes safety issues where LLMs interact with potentially harmful models. BadChain (Xiang et al., 2024) also reveals vulnerabilities of LLMs outputting unintended malicious content under backdoor attack when employing CoT prompting. Chain-of-Feedback (Ahn and Shin, 2024) conducts another evaluation, demonstrating that by repeatedly providing LLMs with non-informative prompts like "*make another attempt*", the quality of responses gradually decreases. These methods underscore the importance of nuanced evaluations of LLMs.

5 Future Directions

While LLMs have demonstrated remarkable abilities in step-by-step problem-solving for various tasks, several challenges remain to be addressed.

Causal Analysis on Intermediates Existing works generally focus on improving task-specific generative results. However, understanding and explaining the underlying mechanisms of LLM reasoning is also essential in realistic scenarios. For example, Wang et al. (2023f) show that LLMs may skip rational steps when generating final results. Wang et al. (2023a) observe a performance gain from CoT even with invalid rationales. These observations indicate the value of a causal analysis on how intermediate steps truly affect the final results.

Reducing Inference Cost A chain leading to the final node of generation often requires multiple sequential inference steps, which are computationally heavy and time-consuming, especially with LLMs. Future research may explore reducing the length of CoX chains while maintaining the quality of generation. It would also be worth studying whether the intermediate nodes of CoX could be executed in parallel or jointly within a single inference step.

Knowledge Distillation The knowledge elicited by the intermediate nodes of CoX contains fine-grained task instructions, which can benefit the training of smaller student models when using a teacher LLM for knowledge distillation. Li et al. (2023b) and Hsieh et al. (2023) have shown that the student model can effectively learn from the rationales of CoT generated by an LLM. Nonetheless, it remains an open question whether the intermediate nodes from broader CoX methods are equally informative in inspiring student learning.

End-to-End Fine-tuning One drawback of CoX is that it does not follow an end-to-end paradigm; i.e., generation errors may accumulate along the chain when self-correction (Le et al., 2024; Dhuliawala et al., 2023) is not enforced. Future research can explore fine-tuning LLMs with CoX prompting and penalizing errors from the final output. By reducing the generation errors end-to-end, we expect this will improve the quality of both the intermediate and final nodes in CoX.

6 Conclusion

This survey explored Chain-of-X methods, building upon the concept of Chain-of-Thought. By categorizing them based on nodes and tasks, we provide a comprehensive overview that highlights the potential of CoX in enhancing LLM capabilities and opens new avenues for future research. Through this survey, we aim to inspire further exploration in a deeper understanding and more creative use of CoX paradigms for LLMs.

Limitations

This survey presents a comprehensive overview of Chain-of-X methods. We have made our best efforts to collect studies leveraging the CoX concept, regardless of whether they are explicitly named as such. However, with the rapidly growing number of works in the field, there is still a chance that some have been missed. We welcome suggestions from the research community and will continue our efforts to survey and update the collected works.

References

Griffin Adams, Alex Fabbri, Faisal Ladhak, Eric Lehman, and Noémie Elhadad. 2023. From sparse to dense: GPT-4 summarization with chain of density prompting. In *Proceedings of the 4th New Frontiers in Summarization Workshop*, pages 68–74, Singapore. Association for Computational Linguistics.

Garima Agrawal, Tharindu Kumarage, Zeyad Alghami, and Huan Liu. 2023. Can knowledge graphs reduce hallucinations in llms?: A survey. *arXiv preprint arXiv:2311.07914*.

Jinwoo Ahn and Kyuseung Shin. 2024. Recursive chain-of-feedback prevents performance degradation from redundant prompting. *arXiv preprint arXiv:2402.02648v2*.

Songlin Bao, Tiantian Li, and Bin Cao. 2024. Chain-of-event prompting for multi-document summarization by large language models. *International Journal of Web Information Systems*.

Maciej Besta, Florim Memedi, Zhenyu Zhang, Robert Gerstenberger, Nils Blach, Piotr Nyczyk, Marcin Copik, Grzegorz Kwaśniewski, Jürgen Müller, Lukas Gianinazzi, et al. 2024. Topologies of reasoning: Demystifying chains, trees, and graphs of thoughts. *arXiv preprint arXiv:2401.14295*.

Rishabh Bhardwaj and Soujanya Poria. 2023. Red-teaming large language models using chain of utterances for safety-alignment. *arXiv preprint arXiv:2308.09662*.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. Chateval: Towards better LLM-based evaluators through multi-agent debate. In *The Twelfth International Conference on Learning Representations*.

Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2023a. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions on Machine Learning Research*.

Zhipeng Chen, Kun Zhou, Beichen Zhang, Zheng Gong, Xin Zhao, and Ji-Rong Wen. 2023b. ChatCoT: Tool-augmented chain-of-thought reasoning on chat-based large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14777–14790, Singapore. Association for Computational Linguistics.

Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2024. Navigate through enigmatic labyrinth a survey of chain of thought reasoning: Advances, frontiers and future. In *The 62nd Annual Meeting of the Association for Computational Linguistics: ACL 2024, Bangkok, Thailand, August 11–16, 2024*. Association for Computational Linguistics.

Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2023. Chain-of-verification reduces hallucination in large language models. *arXiv preprint arXiv:2309.11495*.

Xuan Long Do, Kenji Kawaguchi, Min Yen Kan, and Nancy F Chen. 2023. Choire: Characterizing and predicting human opinions with chain of opinion reasoning. *arXiv preprint arXiv:2311.08385*.

Xi Fang, Weijie Xu, Fiona Anting Tan, Jiani Zhang, Ziqing Hu, Yanjun Qi, Scott Nickleach, Diego Socolinsky, Srinivasan Sengamedu, and Christos Faloutsos. 2024. Large language models(llms) on tabular data: Prediction, generation, and understanding – a survey. *arXiv preprint arXiv:2402.17944*.

Silin Gao, Jane Dwivedi-Yu, Ping Yu, Xiaoqing Ellen Tan, Ramakanth Pasunuru, Olga Golovneva, Koustuv Sinha, Asli Celikyilmaz, Antoine Bosselut, and Tianlu Wang. 2024. Efficient tool use with chain-of-abstraction reasoning. *arXiv preprint arXiv:2401.17464*.

Peiyuan Gong and Jiaxin Mao. 2023. Coascore: Chain-of-aspects prompting for nlg evaluation. *arXiv preprint arXiv:2312.10355*.

Guangzeng Han, Weisi Liu, Xiaolei Huang, and Brian Borsari. 2024. Chain-of-interaction: Enhancing large language models for psychiatric behavior understanding by dyadic contexts. *arXiv preprint arXiv:2403.13786*.

Shirley Anugrah Hayati, Taehee Jung, Tristan Bodding-Long, Sudipta Kar, Abhinav Sethy, Joo-Kyung Kim, and Dongyeop Kang. 2024. Chain-of-instructions: Compositional instruction tuning on large language models. *arXiv preprint arXiv:2402.11532*.

Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8003–8017, Toronto, Canada. Association for Computational Linguistics.

Chenxu Hu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Zhao, and Hang Zhao. 2023a. Chatdb: Augmenting llms with databases as their symbolic memory. *arXiv preprint arXiv:2306.03901*.

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

Hanxu Hu, Hongyuan Lu, Huajian Zhang, Wai Lam, and Yue Zhang. 2023b. Chain-of-symbol prompting elicits planning in large langauge models. *arXiv preprint arXiv:2305.10276*.

Fan Huang, Haewoon Kwak, and Jisun An. 2023. Chain of explanation: New prompting method to generate quality natural language explanation for implicit hate speech. In *Companion Proceedings of the ACM Web Conference 2023*, pages 90–93.

Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards reasoning in large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065, Toronto, Canada. Association for Computational Linguistics.

Qiang Huang, Feng Huang, DeHao Tao, YueTong Zhao, BingKun Wang, and YongFeng Huang. 2024. Coq: An empirical framework for multi-hop question answering empowered by large language models. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 11566–11570. IEEE.

Tatsuro Inaba, Hirokazu Kiyomaru, Fei Cheng, and Sadao Kurohashi. 2023. MultiTool-CoT: GPT-3 can use multiple external tools with chain of thought prompting. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1522–1532, Toronto, Canada. Association for Computational Linguistics.

Bowen Jin, Chulin Xie, Jiawei Zhang, Kashob Kumar Roy, Yu Zhang, Suhang Wang, Yu Meng, and Jiawei Han. 2024. Graph chain-of-thought: Augmenting large language models by reasoning on graphs. *arXiv preprint arXiv:2404.07103*.

Taehee Kim, Yeongjae Cho, Heejun Shin, Yohan Jo, and Dongmyung Shin. 2024. Generalizing visual question answering from synthetic to human-written questions via a chain of qa with a large language model. *arXiv preprint arXiv:2401.06400*.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.

Aditya Kuppa, Nikon Rasumov-Rahe, and Marc Voses. 2023. Chain of reference prompting helps llm to think like a lawyer. In *Generative AI+ Law Workshop*.

Emanuele La Malfa, Christoph Weinhuber, Orazio Torre, Fangru Lin, Anthony Cohn, Nigel Shadbolt, and Michael Wooldridge. 2024. Code simulation challenges for large language models. *arXiv preprint arXiv:2401.09074*.

Hung Le, Hailin Chen, Amrita Saha, Akash Gokul, Doyen Sahoo, and Shafiq Joty. 2024. Codechain: Towards modular code generation through chain of self-revisions with representative sub-modules. In *The Twelfth International Conference on Learning Representations*.

Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. 2023a. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*.

Yoon Kyung Lee, Inju Lee, Minjung Shin, Seoyeon Bae, and Sowon Hahn. 2023b. Chain of empathy: Enhancing empathetic response of large language models based on psychotherapy models. *arXiv preprint arXiv:2311.04915*.

Deren Lei, Yaxi Li, Mingyu Wang, Vincent Yun, Emily Ching, Eslam Kamal, et al. 2023. Chain of natural language inference for reducing large language model ungrounded hallucinations. *arXiv preprint arXiv:2310.03951*.

Chengshu Li, Jacky Liang, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, and Brian Ichter. 2023a. Chain of code: Reasoning with a language model-augmented code emulator. *arXiv preprint arXiv:2312.04474*.

Liunian Harold Li, Jack Hessel, Youngjae Yu, Xiang Ren, Kai-Wei Chang, and Yejin Choi. 2023b. Symbolic chain-of-thought distillation: Small models can also “think” step-by-step. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2665–2679, Toronto, Canada. Association for Computational Linguistics.

Xingxuan Li, Ruochen Zhao, Yew Ken Chia, Bosheng Ding, Shafiq Joty, Soujanya Poria, and Lidong Bing. 2024a. Chain-of-knowledge: Grounding large language models via dynamic knowledge adapting over heterogeneous sources. In *The Twelfth International Conference on Learning Representations*.

Yangning Li, Shirong Ma, Xiaobin Wang, Shen Huang, Chengyue Jiang, Hai-Tao Zheng, Pengjun Xie, Fei Huang, and Yong Jiang. 2024b. Ecomgpt: Instruction-tuning large language models with chain-of-task tasks for e-commerce. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18582–18590.

Zhan Ling, Yunhao Fang, Xuanlin Li, Zhiao Huang, Mingu Lee, Roland Memisevic, and Hao Su. 2023. Deductive verification of chain-of-thought reasoning. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Hanmeng Liu, Zhiyang Teng, Leyang Cui, Chaoli Zhang, Qiji Zhou, and Yue Zhang. 2023. LogiCoT: Logical chain-of-thought instruction tuning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2908–2921, Singapore. Association for Computational Linguistics.

Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. 2024a. Chain of hindsight aligns language models with feedback. In *The Twelfth International Conference on Learning Representations*.

Zhili Liu, Yunhao Gou, Kai Chen, Lanqing Hong, Jiahui Gao, Fei Mi, Yu Zhang, Zhenguo Li, Xin Jiang, Qun Liu, and James T. Kwok. 2024b. Mixture of insightful experts (mote): The synergy of thought chains and expert mixtures in self-alignment. *arXiv preprint arXiv:2405.00557*.

Zuyan Liu, Yuhao Dong, Yongming Rao, Jie Zhou, and Jiwen Lu. 2024c. Chain-of-spot: Interactive reasoning improves large vision-language models. *arXiv preprint arXiv:2403.12966*.

Hongyuan Lu, Haoyang Huang, Dongdong Zhang, Haoran Yang, Wai Lam, and Furu Wei. 2023. Chain-of-dictionary prompting elicits translation in large language models. *arXiv preprint arXiv:2305.06575*.

Ruilin Luo, Tianle Gu, Haoling Li, Junzhe Li, Zicheng Lin, Jiayi Li, and Yujiu Yang. 2024. Chain of history: Learning and forecasting with llms for temporal knowledge graph completion. *arXiv preprint arXiv:2401.06072*.

Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 305–329, Nusa Dua, Bali. Association for Computational Linguistics.

Weiyu Ma, Qirui Mi, Xue Yan, Yuqiao Wu, Runji Lin, Haifeng Zhang, and Jun Wang. 2023. Large language models play starcraft ii: Benchmarks and a chain of summarization approach. *arXiv preprint arXiv:2312.11865*.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. Self-refine: Iterative refinement with self-feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Fanxu Meng, Haotong Yang, Yiding Wang, and Muhan Zhang. 2023. Chain of images for intuitively reasoning. *arXiv preprint arXiv:2311.09241*.

Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. 2024. Compositional chain of thought prompting for large multimodal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Debjyoti Mondal, Suraj Modi, Subhadarshi Panda, Rituraj Singh, and Godawari Sudhakar Rao. 2024. Kamcot: Knowledge augmented multimodal chain-of-thoughts reasoning. In *AAAI Conference on Artificial Intelligence*.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Zhenyu Pan, Haozheng Luo, Manling Li, and Han Liu. 2024. Chain-of-action: Faithful and multimodal question answering through large language models. *arXiv preprint arXiv:2403.17359*.

Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. Measuring and narrowing the compositionality gap in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.

Ji Qi, Ming Ding, Weihan Wang, Yushi Bai, Qingsong Lv, Wenyi Hong, Bin Xu, Lei Hou, Juanzi Li, Yuxiao Dong, et al. 2024. Cogcom: Train large vision-language models diving into details through chain of manipulations. *arXiv preprint arXiv:2402.04236*.

Xihe Qiu, Teqi Hao, Shaojie Shi, Xiaoyu Tan, and Yu-Jie Xiong. 2024. Chain-of-lora: Enhancing the instruction fine-tuning performance of low-rank adaptation on diverse instruction set. *IEEE Signal Processing Letters*.

Sergio Servantez, Joe Barrow, Kristian Hammond, and Rajiv Jain. 2024. Chain of logic: Rule-based reasoning with large language models. *arXiv preprint arXiv:2402.10400*.

Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024. Visual cot: Unleashing chain-of-thought reasoning in multi-modal language models. *arXiv preprint arXiv:2403.16999*.

Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni, Heung-Yeung Shum, and Jian Guo. 2024. Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph. In *The Twelfth International Conference on Learning Representations*.

Mingxu Tao, Dongyan Zhao, and Yansong Feng. 2024. Chain-of-discussion: A multi-model framework for complex evidence-based question answering. *arXiv preprint arXiv:2402.16313*.

Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, Toronto, Canada. Association for Computational Linguistics.

Kohei Uehara, Nabarun Goswami, Hanqin Wang, Toshiaki Baba, Kohtaro Tanaka, Tomohiro Hashimoto, Kai Wang, Rei Ito, Takagi Naoya, Ryo Umagami, et al. 2024. Advancing large multi-modal models with explicit chain-of-reasoning and visual question generation. *arXiv preprint arXiv:2401.10005*.

Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun. 2023a. Towards understanding chain-of-thought prompting: An empirical study of what matters. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2717–2739, Toronto, Canada. Association for Computational Linguistics.

Hanbin Wang, Zhenghao Liu, Shuo Wang, Ganqu Cui, Ning Ding, Zhiyuan Liu, and Ge Yu. 2023b. Intervenor: Prompt the coding ability of large language models with the interactive chain of repairing. *arXiv preprint arXiv:2311.09868*.

Hongru Wang, Rui Wang, Fei Mi, Yang Deng, Zezhong Wang, Bin Liang, Ruifeng Xu, and Kam-Fai Wong. 2023c. Cue-CoT: Chain-of-thought prompting for responding to in-depth dialogue questions with LLMs. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12047–12064, Singapore. Association for Computational Linguistics.

Jianing Wang, Qiushi Sun, Nuo Chen, Xiang Li, and Ming Gao. 2023d. Boosting language models reasoning with chain-of-knowledge prompting. *arXiv preprint arXiv:2306.06427*.

Keheng Wang, Feiyu Duan, Sirui Wang, Peiguang Li, Yunsen Xian, Chuantao Yin, Wenge Rong, and Zhang Xiong. 2023e. Knowledge-driven cot: Exploring faithful reasoning in llms for knowledge-intensive question answering. *arXiv preprint arXiv:2308.13259*.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023f. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.

Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi, Xingshan Zeng, Wenyong Huang, Lifeng Shang, Xin Jiang, and Qun Liu. 2023g. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.

Zihao Wang, Anji Liu, Haowei Lin, Jiaqi Li, Xiaojian Ma, and Yitao Liang. 2024a. Rat: Retrieval augmented thoughts elicit context-aware reasoning in long-horizon generation. *arXiv preprint arXiv:2403.05313*.

Zilong Wang, Hao Zhang, Chun-Liang Li, Julian Martin Eisenschlos, Vincent Perot, Zifeng Wang, Lesly Miculicich, Yasuhisa Fujii, Jingbo Shang, Chen-Yu Lee, and Tomas Pfister. 2024b. Chain-of-table: Evolving tables in the reasoning chain for table understanding. In *The Twelfth International Conference on Learning Representations*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.

Nan Xi, Jingjing Meng, and Junsong Yuan. 2023a. Chain-of-look prompting for verb-centric surgical triplet recognition in endoscopic videos. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5007–5016.

Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023b. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*.

Wenhan Xia, Chengwei Qin, and Elad Hazan. 2024a. Chain of lora: Efficient fine-tuning of language models via residual learning. *arXiv preprint arXiv:2401.04151*.

Yu Xia, Fang Kong, Tong Yu, Liya Guo, Ryan A. Rossi, Sungchul Kim, and Shuai Li. 2024b. Which llm to play? convergence-aware online model selection with time-increasing bandits. In *Proceedings of the ACM on Web Conference 2024*, WWW '24, page 4059–4070, New York, NY, USA. Association for Computing Machinery.

Yu Xia, Xu Liu, Tong Yu, Sungchul Kim, Ryan Rossi, Anup Rao, Tung Mai, and Shuai Li. 2024c. Hallucination diversity-aware active learning for text summarization. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8665–8677, Mexico City, Mexico. Association for Computational Linguistics.

Yu Xia, Tong Yu, Zhankui He, Handong Zhao, Julian McAuley, and Shuai Li. 2024d. Aligning as debiasing: Causality-aware alignment via reinforcement learning with interventional feedback. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4684–4695, Mexico City, Mexico. Association for Computational Linguistics.

Yuwei Xia, Ding Wang, Qiang Liu, Liang Wang, Shu Wu, and Xiaoyu Zhang. 2024e. Enhancing temporal knowledge graph forecasting with large language models via chain-of-history reasoning. *arXiv preprint arXiv:2402.14382*.

Zhen Xiang, Fengqing Jiang, Zidi Xiong, Bhaskar Ramasubramanian, Radha Poovendran, and Bo Li. 2024. Badchain: Backdoor chain-of-thought prompting for large language models. In *The Twelfth International Conference on Learning Representations*.

Zeqi Xiao, Tai Wang, Jingbo Wang, Jinkun Cao, Wenwei Zhang, Bo Dai, Dahua Lin, and Jiangmiao Pang. 2024a. Unified human-scene interaction via prompted chain-of-contacts. In *The Twelfth International Conference on Learning Representations*.

Ziyang Xiao, Dongxiang Zhang, Yangjun Wu, Lilin Xu, Yuan Jessica Wang, Xiongwei Han, Xiaojin Fu, Tao Zhong, Jia Zeng, Mingli Song, and Gang Chen. 2024b. Chain-of-experts: When LLMs meet complex operations research problems. In *The Twelfth International Conference on Learning Representations*.

Shicheng Xu, Liang Pang, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. 2024. Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks. In *Proceedings of the ACM on Web Conference 2024*, pages 1362–1373.

Yutaro Yamada, Khyathi Chandu, Yuchen Lin, Jack Hessel, Ilker Yildirim, and Yejin Choi. 2024. L3go: Language agents with chain-of-3d-thoughts for generating unconventional objects. *arXiv preprint arXiv:2402.09052*.

Tao Yang, Tianyuan Shi, Fanqi Wan, Xiaojun Quan, Qifan Wang, Bingzhe Wu, and Jiaxiang Wu. 2023. PsyCoT: Psychological questionnaire as powerful chain-of-thought for personality detection. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3305–3320, Singapore. Association for Computational Linguistics.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.

Wenhao Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2023a. Chain-of-note: Enhancing robustness in retrieval-augmented language models. *arXiv preprint arXiv:2311.09210*.

Zihan Yu, Liang He, Zhen Wu, Xinyu Dai, and Jiajun Chen. 2023b. Towards better chain-of-thought prompting strategies: A survey. *arXiv preprint arXiv:2310.04959*.

Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Wang Yongji, and Jian-Guang Lou. 2023. Large language models meet NL2Code: A survey. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7443–7464, Toronto, Canada. Association for Computational Linguistics.

Qingkai Zeng, Yuyang Bai, Zhaoxuan Tan, Shangbin Feng, Zhenwen Liang, Zhihan Zhang, and Meng Jiang. 2024. Chain-of-layer: Iteratively prompting large language models for taxonomy induction from limited examples. *arXiv preprint arXiv:2402.07386*.

Liangyu Zha, Junlin Zhou, Liyao Li, Rui Wang, Qingyi Huang, Saisai Yang, Jing Yuan, Changbao Su, Xiang Li, Aofeng Su, et al. 2023. Tablegpt: Towards unifying tables, nature language and commands into one gpt. *arXiv preprint arXiv:2307.08674*.

Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. 2023a. SpeechGPT: Empowering large language models with intrinsic cross-modal conversational abilities. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15757–15773, Singapore. Association for Computational Linguistics.

Dong Zhang, Xin Zhang, Jun Zhan, Shimin Li, Yaqian Zhou, and Xipeng Qiu. 2024a. Speechgpt-gen: Scaling chain-of-information speech generation. *arXiv preprint arXiv:2401.13527*.

Jiwen Zhang, Jihao Wu, Yihua Teng, Minghui Liao, Nuo Xu, Xiao Xiao, Zhongyu Wei, and Duyu Tang. 2024b. Android in the zoo: Chain-of-action-thought for gui agents. *arXiv preprint arXiv:2403.02713*.

Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023b. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023c. Siren's song in the ai ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.

Zhenduo Zhang, Bowen Zhang, and Guang Liu. 2023d. Coie: Chain-of-instruct editing for multi-attribute face manipulation. *arXiv preprint arXiv:2312.07879*.

Zhuosheng Zhang and Aston Zhang. 2024. You only look at screens: Multimodal chain-of-action agents. *arXiv preprint arXiv:2309.11436*.

Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023e. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.

Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, and Lidong Bing. 2023. Verify-and-edit: A knowledge-enhanced chain-of-thought framework. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5823–5840, Toronto, Canada. Association for Computational Linguistics.

Ge Zheng, Bin Yang, Jiajin Tang, Hong-Yu Zhou, and Sibei Yang. 2023. DDCot: Duty-distinct chain-of-thought prompting for multimodal reasoning in language models. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. 2023. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*.

Enshen Zhou, Yiran Qin, Zhenfei Yin, Yuzhou Huang, Ruimao Zhang, Lu Sheng, Yu Qiao, and Jing Shao. 2024. Minedreamer: Learning to follow instructions via chain-of-imagination for simulated-world control. *arXiv preprint arXiv:2403.12037*.

Jin Ziqi and Wei Lu. 2023. Tab-CoT: Zero-shot tabular chain of thought. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10259–10277, Toronto, Canada. Association for Computational Linguistics.

A Taxonomies of Nodes and Tasks

We present in Figure 4 the complete version of Figure 2 on Chain-of-X taxonomies categorized by nodes and tasks.

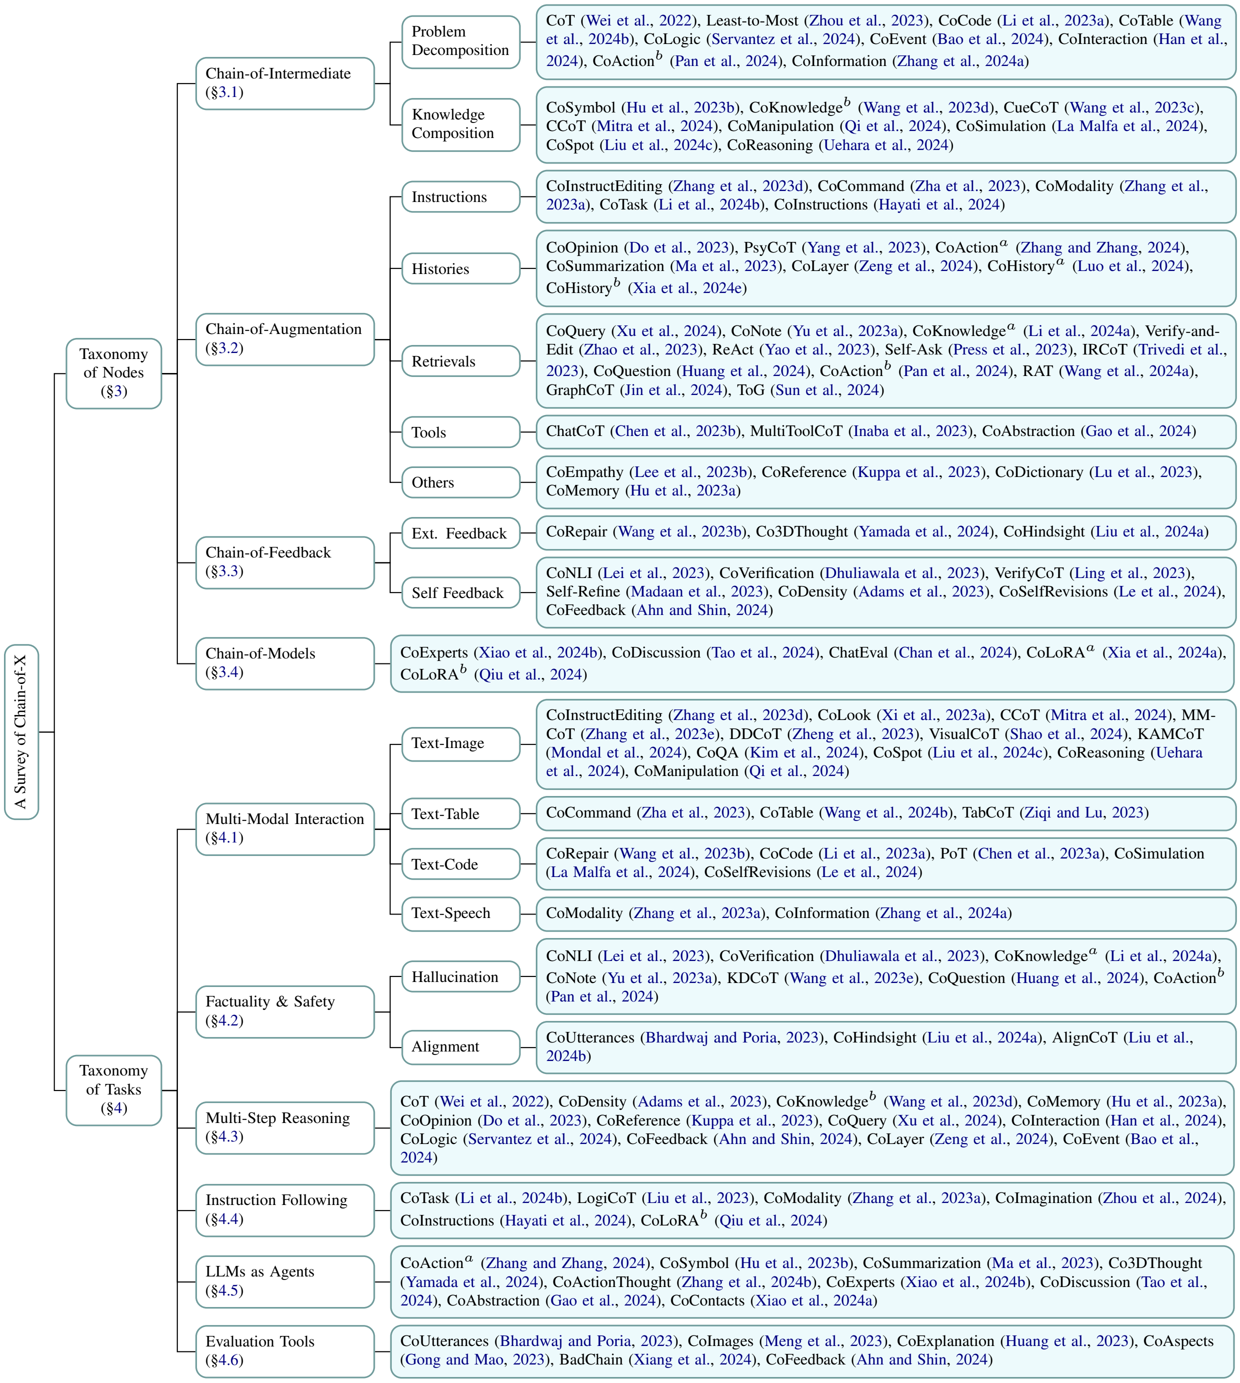

Figure 4: A Survey of Chain-of-X by Taxonomies of Nodes and Tasks.