

ChipMind: Retrieval-Augmented Reasoning for Long-Context Circuit Design Specifications

Changwen Xing^{1,2}, SamZaak Wong^{1,2}, Xinlai Wan², Yanfeng Lu², Mengli Zhang², Zebin Ma², Lei Qi^{2,3}, Zhengxiong Li⁴, Nan Guan⁵, Zhe Jiang^{1,2}, Xi Wang^{*1,2}, Jun Yang^{1,2}

¹School of Integrated Circuits, Southeast University, Nanjing, China

²National Center of Technology Innovation for EDA, Nanjing, China

³School of Computer Science and Engineering, Southeast University, Nanjing, China

⁴Department of Computer Science and Engineering, University of Colorado Denver, USA

⁵Department of Computer Science, City University of Hong Kong, Hong Kong SAR, China

Abstract

While Large Language Models (LLMs) demonstrate immense potential for automating integrated circuit (IC) development, their practical deployment is fundamentally limited by restricted context windows. Existing context-extension methods struggle to achieve effective semantic modeling and thorough multi-hop reasoning over extensive, intricate circuit specifications. To address this, we introduce **ChipMind**, a novel knowledge graph-augmented reasoning framework specifically designed for lengthy IC specifications. ChipMind first transforms circuit specifications into a domain-specific knowledge graph (*ChipKG*) through the *Circuit Semantic-Aware Knowledge Graph Construction* methodology. It then leverages the *ChipKG-Augmented Reasoning* mechanism, combining information-theoretic adaptive retrieval to dynamically trace logical dependencies with intent-aware semantic filtering to prune irrelevant noise, effectively balancing retrieval completeness and precision. Evaluated on an industrial-scale specification reasoning benchmark, ChipMind significantly outperforms state-of-the-art baselines, achieving an average improvement of **34.59%** (up to **72.73%**). Our framework bridges a critical gap between academic research and practical industrial deployment of LLM-aided Hardware Design (LAD).

1 Introduction

As modern integrated circuits (ICs) grow exponentially in scale and complexity, traditional human-driven development has become a critical bottleneck in semiconductor design (International Technology Roadmap for Semiconductors (ITRS) 2001; Foster, Krolnik, and Lacey 2003; Bahr et al. 2020). Large Language Model-Aided Hardware Design (LAD) has emerged as a transformative paradigm (Vaswani et al. 2017; Koroteev 2021), offering new pathways for advancing Electronic Design Automation (EDA) by generating code and testbenches (Bhandari et al. 2024) directly from natural language specifications.

However, a fundamental obstacle hinders the adoption of LAD technology in real-world industrial applications: the




Figure 1: Probability density of specification tokens and design code lines in public LAD benchmarks

limited context window of LLMs (Vaswani et al. 2017). Our analysis reveals that prominent LAD benchmarks (e.g., VerilogEval (Liu et al. 2023), RTLLM2.0 (Liu et al. 2024), VGen (Thakur et al. 2023)) operate on inputs of under 1,000 tokens (Figure 1), significantly smaller than actual industrial specifications—ARM’s AMBA APB Protocol (Arm Ltd. 2023) contains 7.2k tokens, NXP’s I2C-bus specification (NXP 2021) spans 49k tokens, and a complex CPU core like the Xuantie C910 (T-Head Semiconductor Co., Ltd. 2021) demands a staggering 195k tokens. This orders-of-magnitude gap severely restricts current research scalability to realistic, complex industrial scenarios.

To address this challenge, two dominant paradigms exist for extending the context window of LLMs: monolithic long-context models, which leverage techniques like RoPE (Ding et al. 2024) and fine-tuning (Chen et al. 2023); and Retrieval-Augmented Generation (RAG) (Wang et al. 2024) methods. However, both approaches commonly suffer from the “lost-in-the-middle” phenomenon (An et al. 2024): LLMs tend to over-rely on local context while neglecting the overall document structure and cross-module logical connections. This issue is particularly acute in chip design, since the very nature of a chip specification is a tightly-coupled logical chain that spans the entire document. Consequently, any fragmented understanding leads to significant inaccuracies or failures in downstream tasks.

Recently, Knowledge Graph (KG)-based RAG (Pan et al.

*Corresponding author: xi.wang@seu.edu.cn

2024) has been explored to enhance global context perception by LLMs. However, we identify two fundamental limitations when applying these approaches to IC technical documentation: **1) Insufficient Semantic Modeling Capability for Complex IC Documentation:** General-purpose KG construction techniques lack the domain-specific semantic precision required to accurately and comprehensively capture intricate interrelations of entities within IC documents. An incomplete and imprecise KG leads to reduced retrieval accuracy, thereby impairing subsequent reasoning performance. **2) Retrieval Completeness Bottleneck in Multi-Hop Reasoning:** Tracing chip logic often requires following signal chains across modules, demanding complete contextual coverage. However, the commonly used fixed Top- K retrieval mechanism lacks the adaptability to support intermediate reasoning steps. As a result, critical information is often missed, a failure mode that becomes particularly pronounced in complex multi-hop reasoning tasks.

Ultimately, the core bottleneck in LAD has shifted from how to generate code to how to enable LLMs to perform deep comprehension and reasoning over vast specifications.

To overcome this core bottleneck, we introduce **ChipMind**: a knowledge graph-augmented reasoning framework tailored for long IC specifications. ChipMind explicitly parses intricate entity-semantic relationships within chip specifications, reconstructing them into a domain-specific knowledge graph (ChipKG), and employs a *Graph-Augmented Reasoning* mechanism with adaptive retrieval. This enables LLMs to iteratively query ChipKG, emulating human experts to accurately explore and verify deep dependency paths. Comprehensive experiments on a newly introduced benchmark for industrial-scale specification reasoning demonstrate significant and consistent advantages of ChipMind, achieving an average improvement of 34.59% over state-of-the-art baselines, with a maximum performance gain of 72.73% compared to GraphRAG (Edge et al. 2024). Our primary contributions are as follows:

1. We are the first to systematically identify and address the core reasoning bottleneck for LLMs in industrial chip design. To this end, we propose **ChipMind**, a novel framework that integrates a domain-specific knowledge graph with multi-hop reasoning.
2. We design a novel **Circuit Semantic-Aware Knowledge Graph Construction** methodology to build a domain-specific ChipKG. It introduces Circuit Semantic Anchors (CSAs) and the Hierarchical Triple Extraction schema to capture design intent and intricate inter-entity relationships missed by generic methods.
3. We introduce the **ChipKG-Augmented Reasoning** mechanism featuring information-theoretic adaptive retrieval and the CSA-guided semantic filtering layer to enhance the completeness and precision of cross-section dependency tracing.
4. We build **SpecEval-QA**, the first benchmark for industrial-scale specification reasoning, incorporating our newly proposed metric, **Atomic-ROUGE**, robust to paraphrasing and expression variance, thereby aligning

evaluation outcomes more closely with human expert judgment.

2 Related Work

2.1 LLMs in Chip Design and Verification

Large Language Models (LLMs) are increasingly being applied to automate IC development (Chang et al. 2024; Fu et al. 2023; Xu et al. 2025). Seminal works have demonstrated feasibility in generating HDL and HLS code (Pei et al. 2024; Wang et al. 2025a; Wong et al. 2024; Niu et al. 2025; Li et al. 2025a; Ye et al. 2025b; Wan et al. 2025b), creating SystemVerilog Assertions or testbenches (Ye et al. 2025a; Hu et al. 2024), and debugging (Xu et al. 2024; Wang et al. 2025b; Yao et al. 2024; Wan et al. 2025a). However, these methods typically operate on isolated, self-contained code snippets and have limited capability to comprehend the cross-document dependencies inherent in industrial-scale circuit specifications. In contrast, our work is specifically designed to enable reasoning across the entire design documentation.

2.2 Retrieval-Augmented Generation (RAG)

RAG is the standard paradigm for grounding LLMs in external documents, but existing methodologies are ill-suited for the highly relational nature of circuit specifications.

Standard RAG. Conventional RAG, whether based on sparse (e.g., BM25 (Robertson, Zaragoza et al. 2009; Trotman, Puurula, and Burgess 2014)) or dense retrieval (Zhao et al. 2024; Zhan et al. 2021), relies on semantic similarity. This approach is effective for single-hop, fact-based queries but fundamentally fails to model the multi-hop logical dependencies in hardware design. Consequently, it often retrieves a mixture of relevant but functionally incongruous information, derailing the reasoning process.

Knowledge Graph-Augmented RAG (KG-RAG). To imbue retrieval with structural awareness, recent works augment RAG with KGs. However, existing methods still lack the requisite precision for the IC domain. For instance, GraphRAG (Edge et al. 2024; Han et al. 2024) relies on LLM-generated summaries, which risks abstracting away the fine-grained details essential for technical accuracy. To address this loss of granularity, HippoRAG 2 (Gutiérrez et al. 2025) innovatively incorporates raw document chunks as nodes directly into the graph. However, its reliance on generic OpenIE techniques yields entity triples that are too coarse to model the complex relationships between entities in circuit specifications. These collective limitations underscore the need for a new KG-RAG paradigm. To this end, we introduce a framework tailored to the unique semantics of chip design, one that constructs a semantically rich and structurally deep knowledge graph.

2.3 Reasoning-Augmented Frameworks

Frameworks like Chain-of-Thought (CoT) (Wei et al. 2022; Lyu et al. 2023) and Tree-of-Thought (ToT) (Yao et al. 2023a; Long 2023) improve LLM reasoning by decomposing problems into intermediate steps. Integrated approaches such as IRCOT (Trivedi et al. 2023) and ReAct (Yao et al.

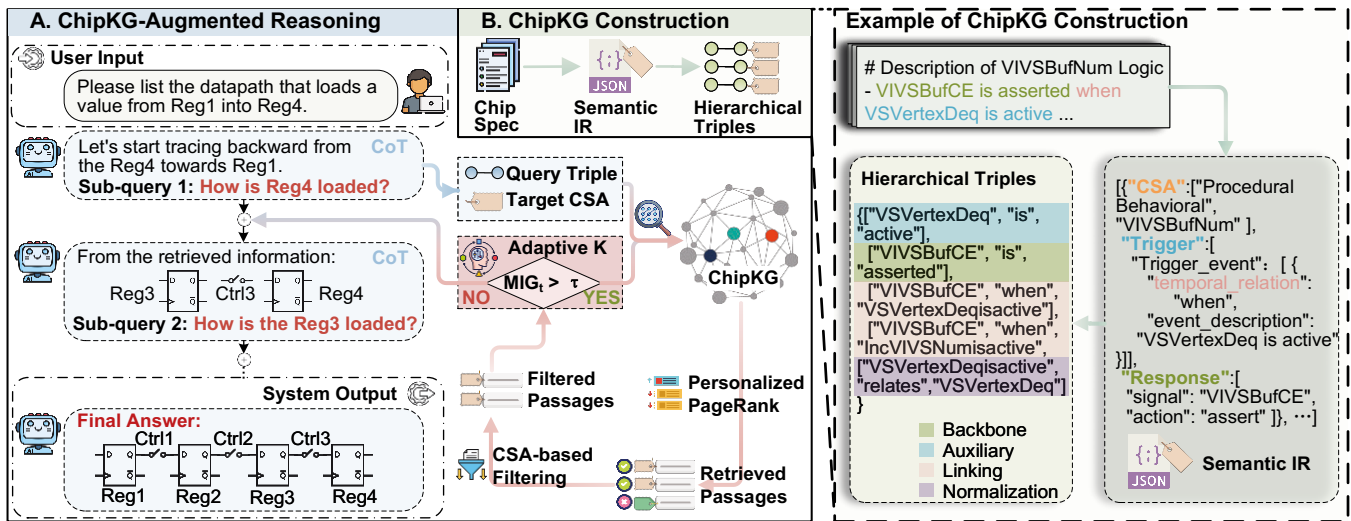


Figure 2: Overview of the ChipMind Framework and an Example of ChipKG Construction

2023b) incorporate retrieval into this process. Their efficacy, however, is fundamentally bottlenecked by the quality of retrieval at each step. Using generic vector-based RAG or web search APIs is ineffective for the specialized and proprietary nature of design documents, as established above. More critically, advanced ToT frameworks like RATT (Zhang et al. 2025) require the reasoning paths to be explicitly predefined. This is incompatible with chip development, where crucial dependencies are often implicit and must be dynamically discovered through iterative exploration. Therefore, our framework departs from these static approaches by introducing a dynamic reasoning process where retrieval and reasoning are synergistically intertwined, allowing the model to actively uncover latent dependencies on-the-fly.

3 Methodology

3.1 ChipMind Framework Overview

To enable deep, multi-hop reasoning over complex Integrated Circuit (IC) design specifications, we propose **ChipMind**, a two-stage framework (shown in Figure 2):

In the first stage, the Circuit Semantic-Aware Knowledge Graph construction methodology is employed. Specifically, semantic content is transformed into a semantic intermediate representation (IR) by functional categorization, and Circuit Semantic Anchors (CSAs) are extracted to identify design intents. Based on semantic IR, structured triples are derived through the Hierarchical Triple Extraction schema, forming a domain-specific Chip Knowledge Graph (ChipKG).

In the second stage, reasoning proceeds through dynamic interactions with ChipKG. When encountering incomplete knowledge during reasoning, the system generates targeted sub-queries, retrieves relevant knowledge graph nodes via vector-similarity matching, and refines their relevance through Personalized PageRank (PPR). Retrieval precision and completeness are further enhanced by employing the adaptive Top-K retrieval strategy based on marginal information gain, along with a CSA-guided semantic filtering

layer to eliminate semantic noise from lexically similar yet irrelevant nodes. This structured integration of knowledge and adaptive reasoning systematically decomposes intricate IC reasoning tasks into transparent, traceable steps.

3.2 Circuit Semantic-Aware Knowledge Graph

Accurate reasoning over chip specifications requires precise modeling of chip design semantics. However, as discussed in the introduction, generic methods fall short in capturing complex, hierarchical semantic relations within IC documents. Therefore, we propose the *Circuit Semantic-Aware Knowledge Graph Construction* methodology, which explicitly deconstructs and reconnects semantic relationships within circuit documents via a tailored paradigm, enabling comprehensive and robust IC domain-specific knowledge graph (ChipKG) construction. A detailed illustration of this methodology is provided in Figure 2.

Circuit Semantic Anchoring and Categorization. We observe that sentences in IC specifications serve two primary functions:

- **Declarative Functional Description:** Describes static properties and definitions, such as module composition, register fields, and signal functions.
- **Procedural Behavioral Description:** Delineates dynamic logic and behavior, such as state transitions, conditional triggers, and signal assignments.

Based on this classification, we utilize specialized parsing templates: declarative descriptions are parsed into central entities and their attributes, while procedural descriptions yield structured “trigger-condition-action” logic. Each sentence is first deconstructed into a concise JSON-based semantic intermediate representation (IR) and distilled into a Circuit Semantic Anchor (CSA). While the IR provides fine-grained details for subsequent triple extraction, the CSA acts as a high-level semantic filter to facilitate efficient downstream reasoning.

Hierarchical Triple Extraction. Based on the extracted IR results, we propose the hierarchical triple extraction schema to comprehensively extract four distinct categories of triples (described below), enabling rich entity relationships within circuit specifications to be restructured into ChipKG.

- **Backbone Triples (T_B):** Encode the central action or definition of an entity.
- **Auxiliary Triples (T_A):** Encode conditional or temporal semantics that qualify the backbone action.
- **Linking Triples (T_L):** Connect T_B with its corresponding T_A s, formally representing the dependency.
- **Normalization Triples (T_N):** Map compound entities to canonical forms, resolving fragmentation and enhancing structural connectivity.

3.3 ChipKG-Augmented Reasoning

To trace complex multi-hop dependencies inherent in chip specifications, we propose *ChipKG-Augmented Reasoning*, enabling the LLM to iteratively query ChipKG. This dynamic interaction is driven by two novel components: (1) Adaptive Top-K Retrieval for comprehensive evidence collection, and (2) Semantic Anchor-Guided Filtering for precise noise reduction. This synergy improves the signal-to-noise ratio for more reliable and interpretable reasoning.

Information-Theoretic Adaptive Top-K Retrieval. As discussed in Section 1, a fundamental challenge in RAG is the static retrieval number, K : a small K misses crucial information, while a large K introduces noise. We thus propose adaptively expanding the retrieval set until marginal utility diminishes.

Formalism. We model the LLM’s belief state as a posterior probability distribution $P(A|C_t)$ over the answer space A , conditioned on context $C_t = (Q, S_t)$, where Q is the query and S_t is the current document set. The value of new information ΔS , is measured by its impact on this belief state. We quantify this impact using the Kullback-Leibler (KL) divergence and define the Marginal Information Gain (MIG) as:

$$MIG(\Delta S | C_t) := D_{KL}(P(A | C_t \cup \Delta S) || P(A | C_t))$$

A significant MIG indicates that ΔS provides novel, decision-relevant evidence, whereas a near-zero value implies redundancy and signals retrieval termination.

Iterative Retrieval Algorithm. As detailed in Algorithm 1, the context set S is iteratively expanded with evidence nodes retrieved from the ChipKG. Since directly computing $P(A|C)$ is infeasible, we approximate Marginal Information Gain (MIG) by prompting the LLM to summarize contexts with (A'_{new}) and without (A'_{base}) new nodes. The MIG at each iteration is then estimated as the cosine distance between their embeddings:

$$MIG_t \approx 1 - \cos(\text{emb}(A'_{base}), \text{emb}(A'_{new}))$$

We terminate the retrieval once MIG_t drops below a threshold τ , indicating diminishing returns.

Algorithm 1: Iterative Context Expansion Algorithm

```

1: Initialize: Form initial set  $S_0$  by retrieving top- $k_0$  documents.
2: for  $t = 0, 1, 2, \dots$  do
3:   Expand: Retrieve next  $\Delta k$  documents to form incremental set  $\Delta S_t$ .
4:   Estimate Belief State (Proxy for intractable  $P(A|C)$ ):
5:     Generate summary  $A'_{base} \leftarrow \text{LLM}(S_t)$ .
6:     Generate summary  $A'_{new} \leftarrow \text{LLM}(S_t \cup \Delta S_t)$ .
7:     Calculate gain  $MIG_t$  based on the change between embeddings of  $A'_{base}$  and  $A'_{new}$ .
8:   Check & Terminate:
9:   if  $MIG_t > \tau$  then
10:     $S_{t+1} \leftarrow S_t \cup \Delta S_t$ . {Information is useful, continue.}
11:   else
12:    return  $S_t$ . {Diminishing returns, terminate.}
13:   end if
14: end for

```

CSA-based Semantic Filtering. While adaptive retrieval ensures evidence completeness, it can also introduce functionally irrelevant information. To maintain retrieval precision, we introduce a semantic filtering layer guided by Circuit Semantic Anchors (CSAs). Specifically, for each query or sub-task Q , we prompt the LLM to infer the query’s intent by generating a target anchor $CSA_{target} = (\text{type}_{target}, \text{entity}_{target})$.

The candidate set S_{cand} is pruned to form the final context S_{final} by retaining only nodes whose anchors exactly match the target anchor CSA_{target} :

$$S_{final} = \{s_i \in S_{cand} \mid \text{Compat}(CSA_i, CSA_{target})\}$$

This precise filtering removes irrelevant information, enhancing signal-to-noise ratio for downstream reasoning.

ChipKG-Augmented Reasoning Workflow. We integrate the preceding components into our ChipKG-Augmented Reasoning framework (Algorithm 2).

An iterative loop begins by reasoning on the current context and introspectively detecting knowledge gaps. Once a gap is identified, ChipMind formulates a targeted sub-query, retrieves relevant evidence from ChipKG, and integrates refined evidence back into reasoning. This loop repeats until sufficient context is collected to synthesize a final, grounded answer.

Our framework advances existing multi-step retrieval methods in two key aspects: we replace static, fixed- K retrieval with a dynamic, information-theoretic strategy to optimize context size, and elevate filtering from shallow semantic matching to intent-level precision using CSA-guided functional alignment. This synergy improves the signal-to-noise ratio for more reliable and interpretable reasoning.

3.4 Atomic-ROUGE

Traditional n-gram overlap metrics (e.g., ROUGE) struggle with semantic variations, while modern metrics like BERTScore improve semantic matching but still fail to reliably evaluate factual correctness—often misled by plausible hallucinations or partially correct statements. To overcome this fundamental limitation, we propose Atomic-ROUGE, a

Algorithm 2: ChipKG-Augmented Reasoning Workflow

Input: An initial query Q_{in} **Output:** A final, grounded answer A_{final}

```
1: Initialize:  $C \leftarrow Q_{in}$ ;  $A_{final} \leftarrow \text{null}$ ;  
2: while termination condition not met do  
3:   // Step 1: Reason and Detect Knowledge Gaps  
4:    $r_t \leftarrow \text{Reason}(C)$ ;  
5:   if DetectUncertainty( $r_t, C$ ) then  
6:     // Step 2: Formulate Information Need  
7:      $(q_t, CSA_{target}) \leftarrow \text{Formulate-Sub-Query}(r_t)$ ;  
8:     // Step 3: Acquire & Refine Knowledge  
9:      $S_{cand} \leftarrow \text{AdaptiveKRetrieve}(q_t, \text{ChipKG})$ ;  
10:     $S_{filtered} \leftarrow \text{CSAGuidedFilter}(S_{cand}, CSA_{target})$ ;  
11:    // Step 4: Integrate and Continue Reasoning  
12:     $C \leftarrow C \cup S_{filtered}$ ;  
13:  else  
14:    // Step 5: Synthesize and Terminate  
15:     $A_{final} \leftarrow \text{Synthesize-Answer}(C)$ ;  
16:    break;  
17:  end if  
18: end while  
19: return  $A_{final}$ 
```

novel metric explicitly designed to evaluate factual fidelity in complex reasoning tasks.

The core of our approach is a two-part decomposition: first, human experts decompose the reference answer y into a set of minimal, self-contained semantic “atomic facts” $A_{ref} = \{a_1, a_2, \dots, a_n\}$. In parallel, a powerful LLM similarly decomposes the generated answer \hat{y} to produce the set A_{gen} . An LLM semantic judge then identifies correct matches $A_{matched}$ via semantic equivalence:

$$A_{matched} = \{\hat{a}_j \in A_{gen} \mid \exists a_i \in A_{ref} \\ \text{s.t. IsSemanticallyEquivalent}(\hat{a}_j, a_i)\}$$

Finally, drawing from classic information retrieval, precision, recall, and F1-score quantify factual correctness, completeness, and balanced semantic fidelity, respectively:

$$P = \frac{|A_{matched}|}{|A_{gen}|}, R = \frac{|A_{matched}|}{|A_{ref}|}, F1 = 2 \cdot \frac{P \cdot R}{P + R}$$

4 Evaluation

To rigorously evaluate ChipMind, we introduce a challenging benchmark specifically designed for industrial-scale Integrated Circuit (IC) specifications. We systematically compare our framework against strong baselines, evaluating performance using our proposed Atomic-ROUGE metric. Furthermore, we also validate the effectiveness and reliability of Atomic-ROUGE itself.

4.1 Experimental Setup

Benchmark. Current benchmarks for LLM-aided hardware design (LAD) fall short of reflecting the scale and complexity of real-world chip specifications. To address this limitation, we propose **SpecEval-QA**, a novel question-answering benchmark derived from a comprehensive, industrial-grade specification of a High-Performance Compute and Interconnect Macro-block (51k tokens).

Question Type	# of Hops	Description
Single-Module Config Loc	1	Identify signals/parameters from single-paragraph text.
Cross-Module Config Loc	5 ~ 12	Locate signals/parameters across multiple paragraphs.
Behavioral Process Analysis	5 ~ 8	Reason about internal module procedures.
Signal Dependency	2 ~ 5	Trace signal flow across modules.
Control Path Tracing	2	Evaluate FSM transitions and logic.

Table 1: Question Types and Descriptions for SpecEval-QA

SpecEval-QA features 25 questions across representative, realistic chip-level reasoning scenarios (Table 1) designed to test long circuit specification comprehension and reasoning over 1~12 reasoning hops. Critically, each question is rigorously annotated with a gold answer, associated atomic facts, and corresponding ground-truth supporting passages. Such detailed annotations enable a fine-grained verification of factual correctness and reasoning pathways, forming the basis for our new metrics.

Baselines. Since no existing end-to-end frameworks directly address this task, we comprehensively evaluate ChipMind against three categories of strong baselines:

- **Vector RAG:** The standard vector-based retrieval methods (BGE-M3 (Chen et al. 2024)) with state-of-the-art (SOTA) LLMs (see Model Selection and Fairness for details), representing the conventional RAG approach.
- **KG-RAG:** Advanced KG-enhanced methods, including GraphRAG (Edge et al. 2024), HippoRAG 2 (Gutiérrez et al. 2025), and LightRAG (Guo et al. 2024), to assess our domain-specific ChipKG against generic graph-based solutions.
- **Reasoning-Augmented Framework:** Advanced reasoning-driven approaches like ReAct (Yao et al. 2023b), IRCoT (Trivedi et al. 2023), and Search-o1 (Li et al. 2025b), specifically benchmarking the effectiveness of ChipMind’s KG-augmented reasoning against contemporary reasoning strategies.

Model Selection and Fairness. We select several top-tier LLMs, including: SOTA closed-source models, GPT-4.1 (OpenAI 2025) and Claude-4 (Anthropic 2025); a strong open-source baseline Llama-4-scout (Meta 2025); and a reasoning-optimized open-source model DeepSeek-R1 (Guo et al. 2025). Because experimental results (Table 2) show that GPT-4.1 achieves the highest performance, we adopt it as the backbone across all KG-RAG and reasoning-augmented frameworks. This standardization ensures observed performance differences reflect framework design rather than LLM capability. The sole exception is Search-o1, whose reasoning process is tightly coupled with its native QwQ-32B-Preview (Qwen Team 2024) model. Addition-

Question Type		Signal-Module Config Loc		Cross-Module Config Loc		Behavioral Process Analysis		Signal Dependency		Control Path Tracing		AVG
Model Type	Model Name	AVG	STD	AVG	STD	AVG	STD	AVG	STD	AVG	STD	
Vector RAG	GPT-4.1	0.82	0.15	0.63	0.26	0.87	0.04	0.85	0.21	0.78	0.17	<u>0.79</u>
	Claude-4	0.76	0.21	0.56	0.38	0.82	0.09	0.82	0.19	0.71	0.31	0.73
	Llama-4-scout	0.93	0.15	0.49	0.25	0.60	0.22	0.60	0.19	0.60	0.32	0.64
	DeepSeek-R1	0.86	0.18	0.60	0.26	0.72	0.13	0.89	0.18	0.64	0.30	0.78
KG-based RAG	GraphRAG	0.56	0.52	0.45	0.48	0.76	0.11	0.40	0.42	0.77	0.17	0.55
	HippoRAG 2	0.94	0.13	0.73	0.25	0.78	0.19	0.68	0.35	0.72	0.18	<u>0.76</u>
	LightRAG	0.68	0.42	0.38	0.15	0.69	0.20	0.63	0.21	0.55	0.18	0.61
Reasoning-augmented Framework	ReAct	0.94	0.13	0.67	0.16	0.73	0.12	0.72	0.30	0.90	0.11	0.79
	IRCoT	0.93	0.15	0.59	0.21	0.88	0.12	0.83	0.18	0.79	0.16	<u>0.81</u>
	Search-o1	0.76	0.43	0.59	0.41	0.68	0.24	0.71	0.29	0.74	0.17	0.70
KG-augmented Reasoning	ChipMind	1.00	0.00	0.84	0.15	0.93	0.02	0.97	0.06	1.00	0.00	0.95

Table 2: Performance Comparison (Atomic-ROUGE F1 Scores) of ChipMind and Baseline Methods

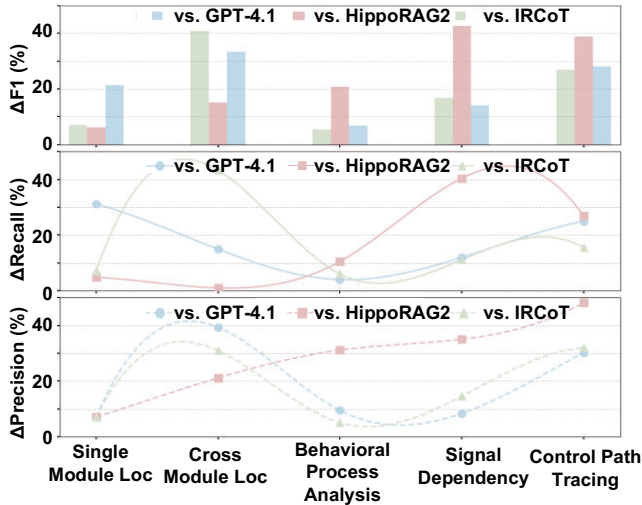


Figure 3: Performance Gains by Task Type

ally, to guarantee fairness in retrieval, all retrieval-dependent frameworks utilize BGE-M3 retriever, accessing an identical data corpus.

Prototyping and Implementation. ChipMind is designed as a model-agnostic framework, compatible with various capable LLMs. In our implementation, we selected DeepSeek-R1 as the core reasoning engine, given its proven effectiveness in logic-intensive tasks and open accessibility that ensures reproducibility. Tasks demanding strong instruction-following capabilities, such as ChipKG construction and Atomic-ROUGE evaluation, leveraged GPT-4.1. We set the sampling temperature to 0.7 for generative tasks and 0.2 for evaluation to improve stability and reproducibility. All experiments were performed on a server with dual Intel Xeon Platinum 8480+ CPUs and eight NVIDIA H20 GPUs (96 GB each).




Figure 4: Retrieval Distribution at Top-K. Top: Half-violin plots illustrating probability density. Bottom: Box-and-whisker plots summarizing distribution characteristics.

4.2 End-to-End Performance

Evaluation Metric. As detailed in Section 3.4, we adopt our proposed Atomic-ROUGE metric for evaluation. By measuring semantic overlap between atomic facts extracted from generated and reference answers, Atomic-ROUGE robustly quantifies precision, recall, and F1-score, effectively handling semantic variations and paraphrasing. We further validate Atomic-ROUGE itself in Section 4.4.

Overall Performance. To ensure robust evaluation, each method was executed 5 times, with each run assessed 20 times using Atomic-ROUGE. Outliers were removed via the 2σ rule before averaging. As summarized in Table 2, ChipMind achieves a SOTA mean F1-score of 0.95, outperforming all baselines by an average of 34.59% and a maximum gain of 72.73% compared to GraphRAG. ChipMind consistently surpasses the strongest baselines within each category, with improvements of 20.25% over GPT-4.1 (F1=0.79), 25% over HippoRAG 2 (F1=0.76), and 17.28% over IRCoT (F1=0.81).

Performance by Task Type. Figure 3 compares ChipMind’s performance improvements across task types against top baselines (GPT-4.1, HippoRAG 2, and IRCoT), highlighting component-level advantages:

For single and cross-module configuration localization

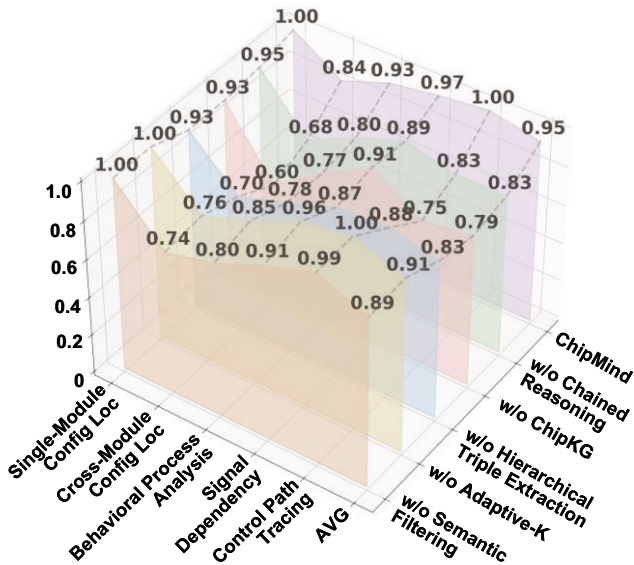


Figure 5: Results of Ablation Experiments

tasks, ChipMind substantially outperforms GPT-4.1 and IR-CoT, which rely on unstructured RAG. This advantage stems from ChipMind’s KG-enhanced retrieval for greater precision, coupled with adaptive iterative retrieval that ensures comprehensive context coverage and improved recall.

For complex reasoning tasks (behavioral processes, signal dependencies, control paths), ChipMind markedly surpasses HippoRAG 2’s single-round static retrieval by decomposing tasks into iterative, verifiable reasoning steps, capturing deep dependencies overlooked by single-step retrieval.

4.3 Retrieval Performance Analysis

To isolate the source of ChipMind’s performance gains, we compare the retrieval completeness of vector-based RAG (BGE-M3), KG-RAG (HippoRAG 2), and ChipMind using **System Recall@ K** , defined as the fraction of ground-truth passages retrieved during the entire reasoning process.

As shown in Fig. 4, ChipMind achieves near-perfect System Recall@20 of 99.2%, significantly outperforming single-pass methods HippoRAG 2 (86.8%) and BGE-M3 (70.5%). This confirms that the baselines’ lower performance primarily stems from incomplete evidence retrieval.

The performance gap arises from semantic mismatches in multi-hop queries, as initial queries are often semantically distant from intermediate evidence, causing single-pass retrieval to rank critical passages poorly. ChipMind bridges this gap via iterative sub-query generation, progressively aligning with intermediate evidence and elevating essential passages to top-ranked positions for complete retrieval.

4.4 Atomic-ROUGE Validation

To validate that Atomic-ROUGE aligns with human judgments, we compared its scores against expert ratings. Three senior chip engineers independently rated ChipMind outputs on 10-point scales for Semantic Fidelity and Answer




Figure 6: Effect of Semantic Filtering vs. Top-K

Completeness; ratings were averaged to form the final human score. Atomic-ROUGE achieved a Pearson correlation of 0.83 with human ratings, surpassing the BERTScore ($r = 0.71$). This confirms Atomic-ROUGE as a reliable automated metric for expert evaluation in fact-intensive tasks.

4.5 Ablation Study

To quantify the contribution of each key component, we conducted four ablation studies, systematically disabling modules from the full ChipMind framework (Figure 5).

Contribution of ChipKG-Augmented Reasoning. Replacing the multi-turn loop with single-pass reasoning sharply decreases performance on multi-hop tasks, confirming the necessity of dynamic sub-query generation for complete evidence retrieval (Section 4.3). Notably, even the degraded single-pass variant (F1=0.83) surpasses KG-RAG baselines such as HippoRAG 2 (0.76), underscoring the inherent schema advantage of our domain-specific ChipKG.

Contribution of ChipKG. Replacing ChipKG with standard Vector RAG severely reduces precision on localization tasks. Further substituting our triple-extraction schema with OpenIE, although outperforming Vector RAG, remains inferior to the full system. This highlights our ChipKG schema’s crucial role in modeling the logic of chip specifications.

Contribution of Adaptive-K Retrieval. The adaptive- K mechanism is the most critical for cross-module configuration localization tasks, as dynamically expanding the search scope ensures comprehensive information coverage.

Contribution of Semantic Filtering. The CSA-guided filter mitigates context pollution, with performance sharply degrading as K increases without it, due to noise from irrelevant documents (Figure 6). In contrast, ChipMind maintains robust performance even at $K = 50$, highlighting the filter’s role in preserving signal-to-noise ratios and supporting larger retrieval budgets.

5 Conclusion

This paper introduced ChipMind, a graph-enhanced reasoning framework designed for the unique complexities of hardware specifications. ChipMind deconstructs implicit logic from specification text into an explicit Chip Knowledge Graph (ChipKG), and its graph-augmented engine enables an LLM to perform verifiable, multi-hop reasoning with high contextual sufficiency and precision. This work sheds light on how structured reasoning can enhance the viability of LLMs for industrial LLM-aided hardware design challenges.

Acknowledgments

This work is supported by the National Key Research and Development Program of China (Grant No. 2024YFB4405600) and the National Natural Science Foundation of China (Grant No. 92464301). We gratefully acknowledge the support from these programs.

References

2021. UM10204: I2C-bus specification and user manual. User Manual UM10204, NXP Semiconductors N.V., Eindhoven, The Netherlands. Accessed: 2025-11-12.
- An, S.; Ma, Z.; Lin, Z.; Zheng, N.; Lou, J.-G.; and Chen, W. 2024. Make your llm fully utilize the context. *Advances in Neural Information Processing Systems*, 37: 62160–62188.
- Anthropic. 2025. Introducing Claude 4.
- Arm Ltd. 2023. AMBA APB Protocol Specification. Technical Report IHI 0024E, Arm Ltd. Non-Confidential; ID022823.
- Bahr, R.; Barrett, C.; Bhagdikar, N.; Carsello, A.; Daly, R.; Donovan, C.; Durst, D.; Fatahalian, K.; Feng, K.; Hanrahan, P.; Hofstee, T.; Horowitz, M.; Huff, D.; Kjolstad, F.; Kong, T.; Liu, Q.; Mann, M.; Melchert, J.; Nayak, A.; Niemetz, A.; Nyengele, G.; Raina, P.; Richardson, S.; Setaluri, R.; Setter, J.; Sreedhar, K.; Strange, M.; Thomas, J.; Tornig, C.; Truong, L.; Tsiskaridze, N.; and Zhang, K. 2020. Creating an Agile Hardware Design Flow. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, 1–6.
- Bhandari, J.; Knechtel, J.; Narayanawamy, R.; Garg, S.; and Karri, R. 2024. Llm-aided testbench generation and bug detection for finite-state machines. *arXiv preprint arXiv:2406.17132*.
- Chang, K.; Chen, Z.; Zhou, Y.; Zhu, W.; Wang, K.; Xu, H.; Li, C.; Wang, M.; Liang, S.; Li, H.; et al. 2024. Natural language is not enough: Benchmarking multi-modal generative AI for Verilog generation. In *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, 1–9.
- Chen, J.; Xiao, S.; Zhang, P.; Luo, K.; Lian, D.; and Liu, Z. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv preprint arXiv:2402.03216*.
- Chen, Y.; Qian, S.; Tang, H.; Lai, X.; Liu, Z.; Han, S.; and Jia, J. 2023. Longlora: Efficient fine-tuning of long-context large language models. *arXiv preprint arXiv:2309.12307*.
- Ding, Y.; Zhang, L. L.; Zhang, C.; Xu, Y.; Shang, N.; Xu, J.; Yang, F.; and Yang, M. 2024. Longrope: Extending llm context window beyond 2 million tokens. *arXiv preprint arXiv:2402.13753*.
- Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Metropolitansky, D.; Ness, R. O.; and Larson, J. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Foster, H. D.; Krolnik, A. C.; and Lacey, D. J. 2003. Assertion-Based Verification. *Proceedings of the IEEE*, 91(1): 136–154.
- Fu, Y.; Zhang, Y.; Yu, Z.; Li, S.; Ye, Z.; Li, C.; Wan, C.; and Lin, Y. C. 2023. Gpt4aigchip: Towards next-generation ai accelerator design automation via large language models. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, 1–9. IEEE.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Guo, Z.; Xia, L.; Yu, Y.; Ao, T.; and Huang, C. 2024. Lightrag: Simple and fast retrieval-augmented generation. *arXiv preprint arXiv:2410.05779*.
- Gutiérrez, B. J.; Shu, Y.; Qi, W.; Zhou, S.; and Su, Y. 2025. From rag to memory: Non-parametric continual learning for large language models. *arXiv preprint arXiv:2502.14802*.
- Han, H.; Wang, Y.; Shomer, H.; Guo, K.; Ding, J.; Lei, Y.; Halappanavar, M.; Rossi, R. A.; Mukherjee, S.; Tang, X.; et al. 2024. Retrieval-augmented generation with graphs (graphrag). *arXiv preprint arXiv:2501.00309*.
- Hu, Y.; Ye, J.; Xu, K.; Sun, J.; Zhang, S.; Jiao, X.; Pan, D.; Zhou, J.; Wang, N.; Shan, W.; Fang, X.; Wang, X.; Guan, N.; and Jiang, Z. 2024. UVLLM: An Automated Universal RTL Verification Framework using LLMs. *arXiv:2411.16238*.
- International Technology Roadmap for Semiconductors (ITRS). 2001. *2001 Edition: Design*. Semiconductor Industry Association. Available: <https://www.semiconductors.org/resources/itrs-archive/>.
- Koroteev, M. V. 2021. BERT: a review of applications in natural language processing and understanding. *arXiv preprint arXiv:2103.11943*.
- Li, R.; Xiong, J.; He, X.; Lv, J.; Zhao, J.; and Wang, X. 2025a. ChatHLS: Towards Systematic Design Automation and Optimization for High-Level Synthesis. *arXiv preprint arXiv:2507.00642*.
- Li, X.; Dong, G.; Jin, J.; Zhang, Y.; Zhou, Y.; Zhu, Y.; Zhang, P.; and Dou, Z. 2025b. Search-ol: Agentic search-enhanced large reasoning models. *arXiv preprint arXiv:2501.05366*.
- Liu, M.; Pinckney, N.; Khailany, B.; and Ren, H. 2023. VerilogEval: Evaluating large language models for verilog code generation. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, 1–8. IEEE.
- Liu, S.; Lu, Y.; Fang, W.; Li, M.; and Xie, Z. 2024. OpenLLM-RTL: Open Dataset and Benchmark for LLM-Aided Design RTL Generation (Invited). In *Proceedings of 2024 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. ACM.
- Long, J. 2023. Large language model guided tree-of-thought. *arXiv preprint arXiv:2305.08291*.
- Lyu, Q.; Havaladar, S.; Stein, A.; Zhang, L.; Rao, D.; Wong, E.; Apidianaki, M.; and Callison-Burch, C. 2023. Faithful chain-of-thought reasoning. In *The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP-AACL 2023)*. Meta. 2025. Llama 4.

- Niu, J.; Liu, X.; Niu, D.; Wang, X.; Jiang, Z.; and Guan, N. 2025. ReChisel: Effective automatic chisel code generation by llm with reflection. *arXiv preprint arXiv:2505.19734*.
- OpenAI. 2025. Introducing GPT-4.1 in the API.
- Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; and Wu, X. 2024. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*, 36(7): 3580–3599.
- Pei, Z.; Zhen, H.-L.; Yuan, M.; Huang, Y.; and Yu, B. 2024. BetterV: Controlled Verilog Generation with Discriminative Guidance. *arXiv:2402.03375*.
- Qwen Team. 2024. QwQ: Pondering the Unknown.
- Robertson, S.; Zaragoza, H.; et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4): 333–389.
- T-Head Semiconductor Co., Ltd. 2021. *XuanTie OpenC910 User Manual*. T-Head Semiconductor Co., Ltd. Licensed under the Apache License, Version 2.0.
- Thakur, S.; Ahmad, B.; Fan, Z.; Pearce, H.; Tan, B.; Karri, R.; Dolan-Gavitt, B.; and Garg, S. 2023. Benchmarking large language models for automated verilog rtl code generation. In *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 1–6. IEEE.
- Trivedi, H.; Balasubramanian, N.; Khot, T.; and Sabharwal, A. 2023. Interleaving Retrieval with Chain-of-Thought Reasoning for Knowledge-Intensive Multi-Step Questions. *arXiv:2212.10509*.
- Trotman, A.; Puurula, A.; and Burgess, B. 2014. Improvements to BM25 and language models examined. In *Proceedings of the 19th Australasian Document Computing Symposium*, 58–65.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wan, G.-W.; Su, S.; Wang, R.; Chen, Q.; Wong, S.-Z.; Xing, M.; Feng, H.; Wang, Y.; Zhu, Y.; Zhang, J.; et al. 2025a. Fixme: Towards end-to-end benchmarking of llm-aided design verification. *arXiv preprint arXiv:2507.04276*.
- Wan, G.-W.; Wang, Y.; Wong, S.-Z.; Xiong, J.; Chen, Q.; Zhang, J.; Zhang, M.; Ni, T.; Xing, M.; Hua, Y.; et al. 2025b. GenBen: A Generative Benchmark for LLM-Aided Design. In *Arxiv*.
- Wang, N.; Yao, B.; Zhou, J.; Hu, Y.; Wang, X.; Jiang, Z.; and Guan, N. 2025a. Large language model for verilog generation with code-structure-guided reinforcement learning. In *2025 IEEE International Conference on LLM-Aided Design (ICLAD)*, 164–170. IEEE.
- Wang, N.; Yao, B.; Zhou, J.; Hu, Y.; Wang, X.; Jiang, Z.; and Guan, N. 2025b. Veridebug: A unified llm for verilog debugging via contrastive embedding and guided correction. In *2025 IEEE International Conference on LLM-Aided Design (ICLAD)*, 61–67. IEEE.
- Wang, X.; Wan, G.-W.; Wong, S.-Z.; Zhang, L.; Liu, T.; Tian, Q.; and Ye, J. 2024. Chatcpu: An agile cpu design and verification platform with llm. In *Proceedings of the 61st ACM/IEEE Design Automation Conference*, 1–6.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wong, S.-Z.; Wan, G.-W.; Liu, D.; and Wang, X. 2024. VGV: Verilog generation using visual capabilities of multi-modal large language models. In *2024 IEEE LLM Aided Design Workshop (LAD)*, 1–5. IEEE.
- Xu, K.; Sun, J.; Hu, Y.; Fang, X.; Shan, W.; Wang, X.; and Jiang, Z. 2024. Meic: Re-thinking rtl debug automation using llms. In *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design*, 1–9.
- Xu, Q.; Stok, L.; Drechsler, R.; Wang, X.; Zhang, G. L.; and Markov, I. L. 2025. Revolution or Hype? Seeking the Limits of Large Models in Hardware Design. *arXiv preprint arXiv:2509.04905*.
- Yao, B.; Wang, N.; Zhou, J.; Wang, X.; Gao, H.; Jiang, Z.; and Guan, N. 2024. Location is key: Leveraging large language model for functional bug localization in verilog. *arXiv preprint arXiv:2409.15186*.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2023a. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36: 11809–11822.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2023b. ReAct: Synergizing Reasoning and Acting in Language Models. *arXiv:2210.03629*.
- Ye, J.; Hu, Y.; Xu, K.; Pan, D.; Chen, Q.; Zhou, J.; Zhao, S.; Fang, X.; Wang, X.; Guan, N.; et al. 2025a. From Concept to Practice: an Automated LLM-aided UVM Machine for RTL Verification. *arXiv preprint arXiv:2504.19959*.
- Ye, J.; Liu, T.; Tian, Q.; Su, S.; Jiang, Z.; and Wang, X. 2025b. ChatModel: Automating Reference Model Design and Verification with LLMs. *arXiv preprint arXiv:2506.15066*.
- Zhan, J.; Mao, J.; Liu, Y.; Guo, J.; Zhang, M.; and Ma, S. 2021. Optimizing dense retrieval model training with hard negatives. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 1503–1512.
- Zhang, J.; Wang, X.; Ren, W.; Jiang, L.; Wang, D.; and Liu, K. 2025. Ratt: A thought structure for coherent and correct llm reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 26733–26741.
- Zhao, W. X.; Liu, J.; Ren, R.; and Wen, J.-R. 2024. Dense text retrieval based on pretrained language models: A survey. *ACM Transactions on Information Systems*, 42(4): 1–60.