

Cortex: Achieving Low-Latency, Cost-Efficient Remote Data Access For LLM via Semantic-Aware Knowledge Caching

Chaoyi Ruan¹, Chao Bi², Kaiwen Zheng³, Ziji Shi¹, Xinyi Wan^{4,1}, and Jialin Li¹

¹National University of Singapore, ²USTC, ³University of Toronto, ⁴Sea AI Lab

Abstract

Large Language Model (LLM) agents tackle data-intensive tasks such as deep research and code generation. However, their effectiveness depends on frequent interactions with knowledge sources across remote clouds or regions. Such interactions can create non-trivial latency and cost bottlenecks. Existing caching solutions focus on exact-match queries, limiting their effectiveness for semantic knowledge reuse.

To address this challenge, we introduce Cortex, a novel cross-region knowledge caching architecture for LLM agents. At its core are two abstractions: *Semantic Element* (SE) and *Semantic Retrieval Index* (Seri). A semantic element captures the semantic embedding representation of an LLM query together with performance-aware metadata such as latency, cost, and staticity. Seri then provides a two-stage retrieval pipeline: a vector similar index with semantic embedding for fast candidate selection and a lightweight LLM-powered *semantic judge* for precise validation. Atop these primitives, Cortex builds a new cache interface that includes a new semantic-aware cache hit definition, a cost-efficient eviction policy, and proactive prefetching. To reduce overhead, Cortex co-locates the smaller LLM judge with the main LLM using adaptive scheduling and resource sharing. Our evaluation demonstrates that Cortex delivers substantial performance improvements without compromising correctness. On representative search workloads, Cortex achieves up to a $3.6\times$ increase in throughput by maintaining cache hit rates of over 85% while preserving accuracy virtually identical to non-cached baselines. Cortex also improves throughput for coding tasks by 20%, showcasing its versatility across diverse agentic workloads.

1 Introduction

LLM agents [20, 28, 66] have emerged as powerful tools that autonomously execute tasks such as deep research assistance [33, 39, 59] and sophisticated code generation [5, 14, 44]. They rely on multi-step reasoning loops that strategically query external knowledge from either private knowledge bases via retrieval-augmented generation (RAG) [37] or cloud search services through API calls [26, 54], and they then syn-

thesize the retrieved information step by step to produce a final solution. Yet, this fundamental dependence on external data retrieval introduces a critical bottleneck to agent performance: high latency and cost of external knowledge access.

Compared to non-agent LLM inference, which relies mainly on GPU computation to generate text, LLM agents frequently invoke external data retrieval services/tools such as web search APIs [43, 58] or RAG backend services [10, 30]. Often, the data source and agent models are located in different data centers or regions and are connected by a wide-area network (WAN). Agent deployment thus faces two major challenges: high monetary cost and cross-region latency due to remote tool calling. For applications [27, 53] handling 5-10 million daily queries, costs are substantial. Given the price of each Google Search API call at \$0.005 [23], the application would incur a monthly costs of \$1.5-4.5 million. Furthermore, accessing data sources in remote regions can lead to 300-500 ms end-to-end latencies [56]. Such high latency can impact user experience and disrupt multi-step workflows.

Prior works have considered applying caching to reduce agent cost. However, existing caching strategies are designed to accelerate LLM inference itself, rather than mitigating the cost and latency of external data retrieval. Semantic prompt caches [9, 24, 57, 65] like GPT-Cache [9] store LLM outputs and reuse responses via prompt similarity to skip LLM generation. However, they compare prompts in the embedding space rather than at the knowledge or tool boundary, and they lack validation to ensure that reused content remains correct for the current context or time. Traditional data storage caches, such as key-value store [60], database [13] or file system [21], store KV objects indexed by *exact keys*. These caches lack the ability to evaluate semantic equivalence between non-matching queries. Transformer KV-caches [22, 36, 41] store token KV states to accelerate model decoding, but their scope is limited to inference computation.

To address this gap, we propose a new paradigm, *semantic-aware remote knowledge caching*, to address latency and cost bottlenecks in cross-region data access for LLM agents. Unlike prior solutions that focus on optimizing LLM infer-

ence computation (e.g., GPTCache) or relying on exact-match queries, our approach leverages LLMs’ language understanding to intelligently cache and retrieve external data. Knowledge caching is built on two core abstractions. We first define a *semantic element* (SE) that encapsulates the agent’s query, tool interactions, and the retrieved response. SE is augmented with performance-aware metadata such as latency, cost, and staticity. We then propose *Semantic Retrieval Index (Seri)*, a two-stage retrieval engine that combines 1) an Approximate Nearest Neighbor (ANN) search for high-recall candidate selection and 2) a lightweight LLM-based *Semantic Judge* for precise validation, ensuring true semantic equivalence.

We present Cortex, a concrete implementation of semantic-aware knowledge caching. Cortex seamlessly integrates the power of LLM semantic understanding into a robust caching architecture, significantly reducing external dependencies and costs associated with remote data access. Crucially, Cortex bridges the gap between the uncertainty of semantic matching and the deterministic requirements of a traditional cache. This is achieved by transforming the ANN and semantic judge output into a reliable, semantic-aware cache hit signal. This procedure ensures that only genuinely relevant and contextually appropriate remote information is served. Furthermore, Cortex equips the semantic index with advanced caching policies, such as a cost-efficient adaptive eviction policy and predictive prefetching, driven by SE metadata. Meanwhile, we propose an efficient Cortex implementation using GPU co-location, where the primary agent LLM and the semantic judge LLM can reside on a single GPU managed by a priority-aware scheduler that protects the agent’s critical latency paths.

Our evaluation demonstrates that Cortex delivers substantial performance gains without compromising accuracy. On representative Zipfian and bursty workloads, Cortex achieves up to a 3.6× throughput improvement over exact-match caching by sustaining cache hit rates of over 85%. Crucially, our accuracy analysis reveals that, while a naive semantic cache suffers a significant drop in correctness, Cortex maintains accuracy virtually identical to the non-cached baselines, proving the indispensable role of the semantic judge. This efficiency extends even to complex coding tasks, where Cortex provides a 20% increase in throughput.

2 Background and Motivation

2.1 Agentic Applications Atop LLM

With the rise of reinforcement learning [15, 18], LLMs are empowered with thinking capability, moving beyond simple text generation to generate tool-use commands/queries in intermediate reasoning steps. The most mature ones are searching agents [32] and coding agents [5].

Agentic LLM. Unlike base LLM models [2], primarily constrained by GPU inference speed, the core of an agentic system is an iterative “think-act-observe” loop that repeatedly accesses remote knowledge. As illustrated in Figure 1a, this process begins with a reasoning phase (①, *think*), where the

Table 1: Example cost information for the commonly used remote data access services.

Company	Google	OpenAI	
Operation	Search API	Web Search Preview	Web Search
Cost (per 1k reqs.)	\$5 [23]	\$10-\$25 [50]	\$10 [50]

agent formulates a plan. This plan then materializes when the LLM generates an *action* (i.e., tool call query) to be dispatched over the model context protocol (MCP [4]) to a remote data source. The tool call can be a search query or an API request. The dedicated *think* and *action* type of LLM outputs are encapsulated by special tags like `<think>` and `<tool>`. After retrieving remote knowledge, the agent integrates it into its working context to inform the next cycle. This loop of external data retrieval is not a one-shot step but a repeated operation at the heart of the agent’s problem-solving process, distinguishing its operational profile from non-agentic LLMs.

The Search-R1 [33] example. To make this workflow concrete, consider the behavior of a search-focused agent as depicted in Figure 1b. When tasked with finding the painter of the Mona Lisa, the agent first articulates its internal goal: `<think> I need to find out who painted the Mona Lisa. </think>`. It then translates this thought into a concrete action, issuing a search tool call with the query *who painted the Mona Lisa?*. This action triggers a remote tool call to an external search engine. Upon receiving the result, the agent observes it in an `<info>` block. Each search and info pair represents a single, costly round trip to an external data source.

2.2 Performance and Cost Implications of Remote Data Retrieval.

Heavy dependence on remote data sources creates significant latency and cost implications. Consider a standard cross-region deployment: an agent powered by a 7B model [33, 34] runs on a single H100 GPU in one region while its external tools (e.g., a Search API) reside in another, introducing 100–300 ms network delay for every call.

Latency. External retrieval often rivals the model’s own inference time. As shown in Figure 1c, running Search-R1 on a single H100, external data retrieval constitutes around 40%-50% of the total execution time. While batching multiple requests can improve GPU utilization, the latency penalty of remote tool call is fundamental: each individual request must still wait for its remote tool calls to complete, as these calls lie on the critical path of the agent’s reasoning loop. This creates a latency floor that cannot be reduced via parallelism alone, directly impacting user-perceived response time.

Financial cost. For large-scale deployments, API fees can become a substantial operational expense. Take Google AI mode that handles 100 million MAU [53] as an example. With an average of 10 tool calls per query [52], the service would incur \$0.005 per call, or roughly \$150,000 in daily API fees, as per Table 1. OpenAI charges up to \$25 per 1000 tool calls for Web Search functionality [50], which is more expen-




Figure 1: Agentic workflow and analysis: (a) shows agent serving architecture which interacts with remote data service; (b) is a specific example from Search-R1 [32]; (c) presents breakdown of search queries when running Search-R1 7B on H100.

sive. To serve such query volume with a production-grade dense model (e.g., Qwen2.5-72B or Llama-3-70B), a deployment requires approximately 50–100 H100 GPUs (at roughly \$1.49/hour each [1], or \$1,800–\$3,600/day). For this service, the \$150,000 daily API cost is 40–80× higher than that of GPU compute, making external data access the dominant operational expense.

API rate limit. Beyond latency and cost, external API rate limits raise practical challenges for smaller-scale deployments. For instance, commercial cloud APIs like Google Cloud Search [23] impose strict rate limits, e.g., 100 queries per minute per user. Large enterprises can circumvent this by maintaining multiple accounts or negotiating higher quotas, but such solutions are impractical for individual developers and small teams. As a result, many resort to third-party aggregation services (e.g., SerpAPI), which add both cost and external dependencies. Cortex offers a lightweight alternative that reduces the number of external calls in the first place.

2.3 Access Pattern of the Agentic Workloads

Agentic LLM workloads heavily rely on remote tool calls, whose cost and latency often dominate performance. Caching can mitigate these overheads only if workloads exhibit sufficient locality. To assess this potential, we analyze two representative domains: AI-assisted search and code generation, both of which show statistical structures favorable to caching.

Search-oriented agent workloads. Agentic search workloads are ultimately shaped by real-world human interests, which already drive global web search and are becoming more prominent as major engines adopt AI-powered modes [12, 42]. Google Search, for instance, now incorporates agentic features—multi-source synthesis, conversational follow-ups, and contextual memory—that resemble the multi-step reasoning of LLM agents. Since commercial query logs are proprietary, we use public Google Trends data as a forward-looking proxy for the query distribution faced by large-scale agentic search systems.

Zipfian Distribution of Search Interests. Search queries follow a Zipfian distribution: a few topics draw most traffic while the majority form a long tail. As shown in Figure 2, head queries




Figure 2: Google Trends data show that the top five topics over different time periods follow a Zipfian pattern. Approximate absolute volumes are used (e.g., “20K+” recorded as 20K), highlighting patterns in AI-assisted search and ranking.




Figure 3: Empirical evidence of search queries resulting in bursty and correlated search patterns. A relevant external event leads to a surge in their search interest, as well as that of the topics with related themes.

File-ID	1	2	3	4	5	6	7	8	9
Access Freq.	1.0	0.28	0.22	0.14	0.1	0.08	0.04	0.04	0.04

Table 2: Access frequency of code file for SQLFluff [11] repo and coding problems are from SWE-Bench [47] Dev.

like *Veronika Slowikowska* and *Graham Greene* dominate 24-hour and 7-day ranges, whereas thousands of others form the long tail. Caching these “head” topics can thus yield high hit rates with modest storage, greatly reducing remote calls.

Bursty and Correlated Query Patterns. Beyond a skewed global distribution, the temporal dynamics of queries reveal a second critical characteristic: they are often bursty and corre-

lated. Interest in specific topics can experience sudden spikes in response to external events. As highlighted in [Figure 3](#), these patterns range from technology breakthrough events (OpenAI releasing GPT-5) to major world news (the passing of Elizabeth II and Charles III’s accession to the throne). The key implication for system design is twofold: caching strategies must be time-adaptive to absorb sudden spikes and can exploit topic correlations for proactive prefetching.

Code-agent workloads. While search-driven agents reflect global information demand, agentic LLMs are also widely used for code generation and software maintenance, which show a similarly skewed access pattern. To examine this, we analyzed coding tasks from SWE-Bench [47] on the `sqlfluff` [11] repository and measured how often each file is needed across tasks. As shown in [Table 2](#), file accesses follow a near-Zipfian distribution: one file is required by nearly all tasks, a few core modules are reused heavily, and most files are rarely touched. This long-tail pattern parallels the popularity skew seen in search workloads. This also implies that caching these hot files locally can eliminate many redundant cross-region fetches, and semantic matching can further capture requests that refer to the same file.

2.4 Motivation and Cache Limitations

Motivation. Search-style and coding tasks show Zipfian popularity with bursty reuse, suggesting that caching could greatly reduce redundant remote calls. Yet, user queries in agentic workloads are semantically aligned rather than exactly key-matched, and they run in dynamic environments with complexities like cross-region latency, retrieval cost, and tool API limits. A useful cache must handle *semantic matching* among queries while accounting for *cost*, *cross-region latency*, and *tool rate limits*.

Existing caches and limitations. There are three common flavors of caching approaches, as summarized in [Table 3](#), namely, transformer key-value (KV) caches [22, 36, 41], semantic prompt caches [9, 24, 57, 65], and caches for traditional storage [13, 60] or file system [21]. Despite their popularity, existing caching mechanisms exhibit significant drawbacks when handling agentic workloads.

The first type, transformer KV caches, accelerates token generation by reusing intermediate tensors for exactly-matched prompts, but its scope is confined to inference computation and does not extend to the remote data access layer. The second type, semantic prompt caches (e.g., ContextCache [65], VectorQ [57]), attempts to bypass inference entirely by matching user inputs via vector similarity. However, this approach suffers from a precision-recall trade-off: vector proximity often fails to capture true semantic equivalence. For example, VectorQ, despite applying adaptive thresholds for similar prompts, still relies on the limited semantic awareness of vector similarity, so a query like “apple nutrition facts” might incorrectly match “Apple stock price analysis”. Moreover, these approaches are designed for plain LLM inference

rather than agentic workloads. They cache model outputs at the prompt level but do not intercept or optimize external tool calls. For instance, ContextCache targets conversational LLMs and is not designed to handle the tool-calling patterns. Lastly, studies have explored the data storage cache [51] in databases or file systems in cross-region use cases. However, their exact-match mechanisms are fundamentally inadequate for natural language queries in agent workloads. These systems treat semantically equivalent queries as distinct keys, so even minor wording changes trigger cache misses and unnecessary remote calls. Additionally, none of these caches fully addresses the rate limit impact in the agent scenario, so agent throughput can still be throttled even when cache hits.

Takeaway. Agentic workloads require a cache that performs semantic match at the knowledge or tool boundary, validates correctness, is cost-aware, cross-region, and rate-limit aware so that cache hits yield real end-to-end gains. In contrast to these existing families, Cortex satisfies all required properties as shown in [Table 3](#), motivating its design in the next section.

3 Design Overview

3.1 System Model and Goal

Deployment model. For LLM-based agents that rely on external knowledge and tools, frequent data retrieval from remote sources introduces significant latency and financial cost. Geographic separation between inference and data sources is often inevitable: organizations calling external LLM APIs cannot control where providers deploy their models, and even with local inference, agents frequently invoke external tools (e.g., web search, databases, third-party APIs) hosted in remote regions. In practice, agentic applications typically rely on the MCP protocol [4] or other general RPC [25] over wide-area networks to fetch required data, amplifying these performance and cost penalties. To mitigate these overheads, we introduce Cortex, a novel cross-region caching system designed to store and serve frequently accessed tool results from the perspective of agentic system efficiency.

Optimization Goal. The goal of Cortex is to reduce costly remote data accesses by reusing semantically equivalent knowledge. Instead of relying on literal text similarity or exact matching, Cortex leverages LLM-based semantic matching to identify and serve cached results for new queries with the same intent. This reduces query latency and significantly cuts API or tool usage costs, improving performance and operational efficiency in geo-distributed LLM agent deployments.

3.2 Design Opportunities and Challenges

Strawman design. As discussed in [§2.4](#), semantic prompt caches reuse model outputs based on embedding similarity. A natural strawman design is to extend this idea: embed each cached retrieval query, place all keys in an ANN index, and on a new query, return the value attached to its nearest neighbor. ANN methods [31, 35] can search large embedding sets efficiently using graph structures or quantization. However, this

Table 3: Cache comparison along critical requirements of agentic workloads.

System	Cached data	Semantic Match	Cost-aware Eviction	Cross-region Aware	Rate-limit aware
Transformer KV-cache [22, 36, 41]	Token KV states	✗	✗	✗	✗
Semantic prompt cache [9, 24, 57, 65]	LLM response	✗	✗	✗	✗
Traditional storage cache [13, 21, 60]	KV object	✗	✓	✓	✗
Cortex	External knowledge	✓	✓	✓	✓

naive design exposes key gaps. It treats similarity search as if it were a cache, yet it offers no guarantee that top-ranked items are actually correct, fresh, or cost-efficient to reuse. These issues motivate the following design challenges.

Challenge 1: Ensuring accurate and valid semantic matches. Embeddings with ANN can reliably find and return textually similar queries, but textual or surface-level similarity [16] does not imply semantic equivalence. Two queries can be near in the vector space yet require different answers or reflect different intents. Correctness also depends on context and time. A match for a historical fact may remain valid, whereas a match for a fast-evolving topic can become stale. A practical system needs a validation mechanism that turns noisy similarity scores into trustworthy hit-or-miss decisions while accounting for both contextual fit and temporal validity.

Challenge 2: From similarity search to a real cache. An embedding model and an ANN index are retrieval components, not a cache. They are typically used to search an entire corpus and return a ranked list of candidates. A cache, in contrast, is a capacity-limited, online component that must emit a deterministic hit signal, admit new items, evict old ones, respect freshness, and avoid polluting itself with near-misses. Using ANN as the front door of a cache forces several questions to be answered explicitly: when does a candidate become a cache hit rather than just a similar item; what metadata beyond the embedding are required to reason about staticity, cost, and scope; how should admission and eviction operate when items differ in lifetime, retrieval cost, and size; and how can we prefetch likely next items without undermining precision. Bridging this gap is essential to turn similarity search into a functional cache rather than a best-effort lookup.

Challenge 3: Achieving sophistication with efficiency. Incorporating semantic validation, freshness reasoning, and predictive prefetching improves quality at the expense of higher compute and memory consumption. The system must integrate these capabilities while ensuring that the end-to-end savings from avoided remote calls outweigh local overheads, thus necessitating lightweight models, careful resource sharing with the serving LLM, and scheduling that protects user-facing latency.

3.3 Cortex Architecture Overview

To address these three challenges, Cortex adopts a layered architecture shown in Figure 4. Each component is explicitly designed to tackle a subset of these challenges while working together with other components as an integrated system. The

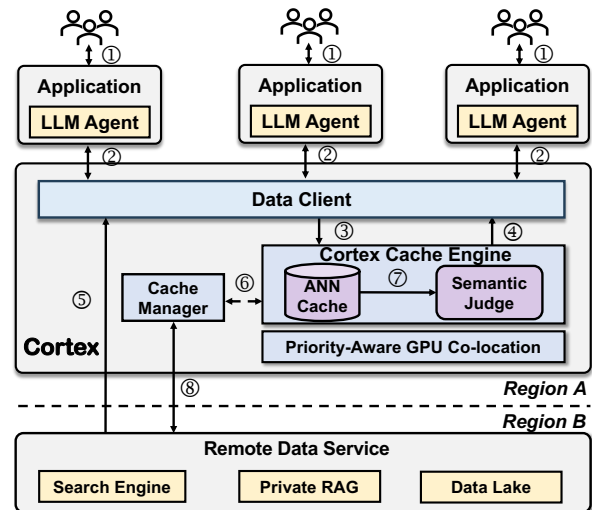


Figure 4: Cortex’s agentic caching architecture where agent LLM and data source reside in two geo-distributed regions

top tier is the *user agentic application*, where an LLM agent generates queries. The bottom tier is the Remote Data Service, the source of data, located in a remote region or cloud with a data retrieval latency around 300 ms-500 ms, following the common industrial practice [40]. This latency encompasses not only network RTT but also API processing and request queuing. Between these two lies the Cortex Engine, which provides a transparent caching interface. It consists of three primary internal components:

Data client. The *data client* serves as the transparent entry point into the Cortex system, which intercepts all outgoing requests generated by the agent before they can reach external services. This interception allows Cortex to seamlessly redirect the agent’s queries into its internal caching workflow without modifying the agent application.

Cortex cache engine. At the core of the engine is the retrieval pipeline, engineered to achieve high-precision semantic matching. This pipeline operates in two distinct stages to balance speed with accuracy. Initially, an ANN index [31] (Faiss [17] is used in our experiments.) based cache rapidly identifies a set of candidate data items based on vector similarity. These candidates are then passed to a lightweight *semantic judge* model (typically small with ~1B parameters), which acts as a validation and is prompted to scrutinize each one for true contextual and semantic relevance to eliminate the false positives that plague naive semantic query caches.

Agent-aware cache manager. Finally, the agent-aware cache

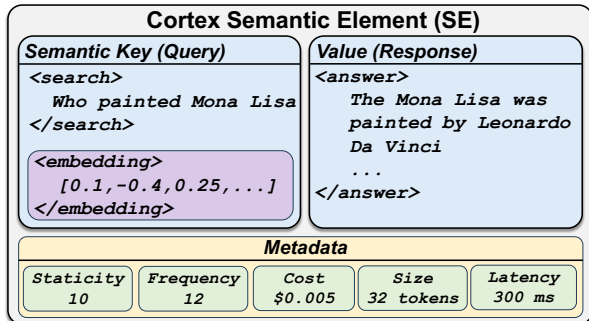


Figure 5: Cortex’s semantic element structure

manager provides the system’s long-term intelligence by intelligently governing the cache’s contents. Moving beyond simple heuristics, it employs two sophisticated strategies. Its adaptive eviction policy uses a utility-based model to decide which items to discard by considering factors like retrieval cost and importance to the agent’s workflow. Complementing this, its proactive prefetching mechanism analyzes historical access patterns to proactively fetch data that is likely to be needed, ensuring the cache is dynamically optimized to hold the most valuable information.

Workflow The workflow begins when the data client transparently intercepts a structured request and its surrounding context generated by the LLM agent (①-②). Instead of immediately dispatching the request externally, it first queries the Cortex cache engine (③). This query initiates a two-stage retrieval process designed for speed and accuracy. First, an ANN index rapidly surfaces a small set of candidate data items. Then, these candidates are scrutinized by the semantic judge (⑦), ensuring accuracy. If the judge confirms a valid match, the cached result is immediately returned to the agent (④). Otherwise, the request proceeds to the appropriate Remote Data Service (⑤). The retrieved response is then returned to the agent and simultaneously stored in the cache as a new Semantic Element for future usage. Concurrently, the agent-aware cache manager observes the sequence of validated queries (⑥), using this history to proactively prefetch data items (⑧) and perform eviction, continuously optimizing the cache’s contents in the background.

4 Cortex’s Design

This section details Cortex’s architecture for transforming agent interactions into reusable knowledge units with efficient caching, retrieval, and scheduling.

4.1 Semantic Element

We begin with the *Semantic Element (SE)*, which is Cortex’s core caching unit. An SE organizes an agent’s *query* and the *retrieval result* for matching, validation, and reuse. As shown in [Figure 5](#), an SE is a key–value pair: the agent’s query or tool action is the semantic key, and the retrieved information is the value. This structure leverages well-formed agentic outputs that wrap steps with tool tags. For example,

the query “Who painted the Mona Lisa” within a `<search>` tag becomes the key, while the snippet in the `<info>` tag is the value. By reliably parsing these tagged blocks, Cortex encapsulates discrete interactions (web searches, API calls, database lookups) into coherent SEs for caching.

An SE also carries metadata essential for lifecycle decisions. First, an embedding model (e.g., Qwen3-0.6B [61]) converts the semantic key into a high-dimensional vector used as a semantic fingerprint and is used for future matching. Concurrently, the semantic judge model (§4.2) estimates a *staticity* score that quantifies the expected validity duration of the query–result pair on a scale of 1–10: higher scores (9–10) indicate stable facts that are unlikely to change, while lower scores (1–3) indicate time-sensitive information. For instance, “Who painted the Mona Lisa?” is *stable* (10), “Who is the current US President?” is *moderate* (5), whereas “Today’s weather in Paris” is *ephemeral* (1). This score impacts the Time-To-Live assignment and the eviction priority in §4.3.

4.2 Seri: Semantic Retrieval Index

Cortex retrieves semantic matches via the *Semantic Retrieval Index (Seri)*. Building on SEs, Cortex implements a high-throughput, dual-stage retrieval pipeline engineered to resolve the fundamental trade-off between retrieval speed and semantic precision. This system is a sophisticated interplay of multiple specialized models, and its core is a formal constrained optimization problem: to minimize end-to-end query latency while rigorously maintaining semantic accuracy.

To formalize this, we define the components and their associated latencies. The primary model is the agentic LLM, a powerful reasoning model that orchestrates tool use; its inference time is denoted as L_{Agent} . The cache itself relies on two smaller, highly optimized models: a lightweight *Embedding LLM* for generating query vectors and a lightweight *semantic judge model* (LSM), a classifier that validates whether a cached result answers a new query and estimates its staticity. The combined latency of a full two-stage cache lookup/check is $L_{\text{CacheCheck}} = L_{\text{ANN}} + L_{\text{LSM}}$. On a cache miss, an external tool call adds L_{Tool} .

Coarse-grained filter. The retrieval process begins with a high-recall candidate selection stage, where the threshold τ_{sim} (e.g., 0.9 in our setup) directly controls the scope of the search. When the Agent LLM generates a new query, q , the system consults an ANN index. A candidate c_i is selected if its similarity satisfies $\text{CosineSimilarity}(q, q_{\text{cached},i}) \geq \tau_{\text{sim}}$, where $q_{\text{cached},i}$ denotes the query of c_i . A lower, more permissive τ_{sim} increases recall, ensuring more potentially relevant items are passed to the next stage, but at the cost of increasing the validation workload. A higher, stricter τ_{sim} reduces this workload but risks prematurely discarding a correct match, thus lowering the potential hit rate.

Fine-grained validation. Selected candidates proceed to a high-precision semantic validation stage, where the LSM’s threshold, τ_{lsm} , is the primary lever for controlling accuracy

(e.g., 0.9 in our experiment). The LSM evaluates whether the cached result $r_{\text{cached},i}$ is a sufficient answer for the new query q , producing a confidence score S_{lsm} . This guards against false positives from surface-level similarity, e.g., “Apple nutrition facts” vs. “Apple stock price” share keywords but require different answers. The LSM rejects such matches by reasoning about semantic equivalence. A cache hit is confirmed only if $S_{\text{lsm}} \geq \tau_{\text{lsm}}$. A higher τ_{lsm} enforces a stricter standard for equivalence, which increases the cache’s precision (fewer false positives) but can decrease the hit rate by rejecting marginally correct answers. A lower τ_{lsm} boosts the hit rate but at the direct risk of serving more incorrect results, thereby lowering precision.

Optimization goal. The system minimizes the expected latency subject to a target precision that is verified *periodically offline*. A hit incurs $L_{\text{hit}} = L_{\text{Agent}} + L_{\text{CacheCheck}}$, while a miss incurs $L_{\text{miss}} = L_{\text{Agent}} + L_{\text{CacheCheck}} + L_{\text{Tool}}$:

$$\underset{\tau_{\text{sim}}, \tau_{\text{lsm}}}{\text{minimize}} \quad E[L] = P_{\text{hit}} \cdot L_{\text{hit}} + (1 - P_{\text{hit}}) \cdot L_{\text{miss}}$$

The formula reflects an accuracy-throughput trade-off: lower τ_{lsm} improves hit rate but risks incorrect results; higher τ_{lsm} preserves accuracy but increases remote call frequency.

Recalibration A fixed τ_{lsm} is brittle under workload drift or changing correctness requirements. Cortex therefore performs periodic offline recalibration to adjust the LSM decision boundary while keeping the agent’s latency-critical path unaffected. This process, detailed in [Algorithm 1](#), begins by creating a high-quality annotated dataset from a sample of recent queries ([line 1-6](#)). Ground truth is obtained by re-issuing sampled queries to the live tool and comparing tool responses against the cached results. The objective is then to find a new threshold, τ'_{lsm} , that aligns the cache’s performance on this dataset with a pre-defined metric, P_{target} . This target precision represents the desired quality standard (e.g., 0.99 for 99% correctness). To find the new threshold, the system calculates a precision curve for the current LSM ([line 7-8](#)) and finds the value that satisfies P_{target} ([line 9](#)), which is then deployed to the live system. Recalibration is decoupled from the request-serving path and does not block agent requests.

Crucially, this recalibration is highly cost-efficient: in our deployment, it samples 5 recent queries per minute, adding negligible overhead. The annotated set can also fine-tune the LSM to better capture subtle semantic differences over time.

4.3 Building Cache Architecture Atop Seri

Although the Seri index can identify semantically matched items, it is not by itself a cache. To bridge the gap, Cortex layers standard cache abstractions on top of the Seri, defining *cache hit*, *cache eviction* and *cache prefetching* and translate semantic matching signals into cache behavior.

Semantic-aware cache hit. Unlike traditional caches where a hit is a simple key lookup, in our system a cached SE is only considered a “hit” after passing the full validation pipeline detailed in [§4.2](#). On a new query, ANN retrieves a small

Algorithm 1: Periodic Threshold Recalibration

Data: Current LSM J_{lsm} , Target Precision P_{target} , Recent Eval Log L_{recent} , Validation Set D_{val}

Result: Recalibrated Threshold τ'_{lsm}

- 1 $D_{\text{sample}} \leftarrow$ Sample diverse subset from L_{recent}
- 2 $D_{\text{annotated}} \leftarrow \emptyset$
- 3 **for** $(q, r_{\text{cached}}, -, -) \in D_{\text{sample}}$ **do**
- 4 $r_{\text{ground}} \leftarrow$ FetchGT(q)
- 5 label \leftarrow EvaluateGT($(q, r_{\text{cached}}), r_{\text{ground}}$)
- 6 $D_{\text{annotated}}.$ Append((q, r_{cached}) , label)
- 7 scores \leftarrow PredictScores($J_{\text{lsm}}, D_{\text{val}}$)
- 8 precision_levels \leftarrow CalcPrecisionCurve(scores)
- 9 $\tau'_{\text{lsm}} \leftarrow$ FindThreshold(precision_levels, P_{target})
- 10 UpdateSystem(τ'_{lsm})
- 11 **return** τ'_{lsm}

Algorithm 2: LCFU Eviction Policy

Data: Cache, capacity

- 1 **Function** CalScore(se)
- 2 ttl \leftarrow se.ExpirationTime - CurrentTime()
- 3 **if** $se.$ Size == 0 or ttl <= 0 **then**
- 4 **return** 0;
- 5 **return** score \leftarrow
- 6 $\frac{\log(se.Freq+1) \times \log(se.Cost \times 10^3 + 1) \times \log(se.Lat+1) \times \log(se.Stat+1)}{se.Size}$
- 6 Cache.RemoveExpired() // TTL purge first
- 7 **if** Cache.Usage() > capacity **then**
- 8 **foreach** se **in** Cache **do**
- 9 se.score \leftarrow CalScore(se)
- 10 **while** Cache.Usage() > capacity **do**
- 11 victim \leftarrow items.PopFirst()
- 12 Cache.Remove(victim)

set of candidates from the cache. Then the LSM (semantic judge) validates equivalence. Only a valid match registers as a cache hit, and that increments an SE’s frequency count. This process effectively transforms a probabilistic similarity score into a definitive, binary hit/miss event, providing a stable, deterministic bedrock for a caching system.

LCFU eviction policy. With a reliable access signal established, we employ a tailored *Least Cost-Efficient and Frequently Used* (LCFU) eviction policy to manage the cache’s contents. The insight behind LCFU is that not all data are equally beneficial for cost and latency savings. Traditional LRU and LFU policies, which prioritize recency and frequency respectively, fail to account for the varying value of cached items. LCFU addresses this by assigning each SE a value_score that quantifies the benefit of retaining it.

As detailed in [Algorithm 2](#), this value_score is a composite metric derived from retrieval latency and cost, frequency, and staticity. Each factor captures a distinct

retention benefit: frequency reflects reuse likelihood, cost and latency reflect savings per hit, and staticity reflects expected validity duration. Then the score is carefully normalized (line 5). It adjusts the raw metrics for frequency and cost to ensure they always contribute positively and meaningfully. This prevents new items or those with low costs from being unfairly penalized, given the fact that the cost per request is less than 1, and taking its logarithm will return a negative value. This combined score is then normalized by the item’s size to provides a direct rationale: keep items that save the most time/money per byte. For instance, data with a high Retrieval Cost and high access Frequency will naturally receive a higher `value_score`, making them less likely to be evicted. Conversely, ephemeral data with a lower Staticity will easily lead to lowering `value_score`, making it a more likely candidate for eviction even if its frequency is high. For one Stable data, with a large `staticity` and high retrieval cost, it will maintain a higher `value_score` even with less frequent access. This multi-attribute approach prevents transient but popular data from displacing enduring content, optimizing the cache for sustained performance and cost savings. When cache capacity is exceeded, items with the lowest `value_score` are evicted, preserving the most valuable content. While a high `value_score` indicates valuable content, such entries can become stale if kept indefinitely. To prevent outdated information from persisting, Cortex integrates an aging mechanism using a user-defined TTL. Each cache entry is assigned a maximum lifespan, after which it is evicted regardless of its `value_score`. This ensures that even high-cost or frequently accessed items are periodically refreshed, maintaining the cache’s correctness while preserving the benefits of value-based retention.

Predictive prefetching. To complement the reactive LCFU eviction policy, our framework employs proactive history-based prefetching to reduce miss latency. Using a lightweight first-order Markov model trained on confirmed cache hits, the system learns query-to-query transition patterns and calculates $P(q_{i+1}|q_i)$. When this probability exceeds a confidence threshold, the predicted query is asynchronously fetched and added to the cache. Prefetched items enter with zero frequency, giving them minimal `value_score`. If later requested and validated, their frequency increments and they become retained members. If unused, their low score makes them prime eviction candidates. This self-correcting design minimizes cache pollution from speculative fetches.

4.4 Resource-Efficient Model Co-location

Although the two-stage retrieval computation is lightweight, the existing serving framework [63] adopts one model per GPU mode. To keep overall latency low and achieve high-precision without doubling hardware costs, Cortex co-locates the agent (e.g. $\sim 7B$) and judge (e.g. $\sim 1B$) efficiently on a single GPU. The challenge is protecting agent latency from judge processing. Our co-location architecture (Figure 6) is

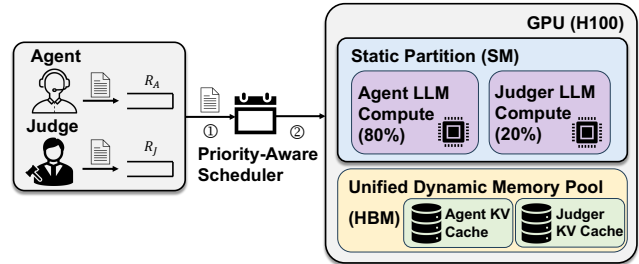


Figure 6: Cortex’s Priority-Aware Scheduling Procedure.

built on a key insight: agent and judge workloads have fundamentally different priorities and resource profiles. Agent work is latency-critical, while judge work is a deferrable optimization. This is because a delayed validation does not block the user; at worst, it is treated as a cache miss, where performance for that single request degrades to the non-cached baseline of a full remote data call. Furthermore, the judge has a smaller memory demand since its prefill-only inference pattern (single token generation), results in a minimal and predictable KV cache footprint. Atop these principles, Cortex builds a two-level defense, to ensure robust and efficient co-location.

Coarse-grained asymmetric partitioning. The foundation of our co-location strategy is to exploit the fundamentally asymmetric nature of the agent and judge workloads. The judge’s workload profile is highly economical: unlike the agent which generates variable-length responses, the judge performs a classification task that yields a single token. This makes its primary consumer of GPU memory (i.e., the KV cache) minimal and highly predictable. Furthermore, its computational needs are lighter than the agent’s complex, multi-step reasoning. To capitalize on this, we use the CUDA Multi-Process Service (MPS) [48] to create a static, asymmetric compute partition, as shown in Figure 6. For instance, we allocate the dominant share of compute resources (e.g., 80%) to the agent and a smaller, sufficient share (e.g., 20%) to the judge. This coarse-grained partitioning optimizes for the common case, maximizing the performance of the primary agent LLM by dedicating the vast majority of resources to it.

Fine-grained prioritization as a high-load guardrail. While static partitioning is effective for the common case, a second, dynamic layer of defense is required to handle periods of high contention and guarantee the agent’s low latency. This is the role of our fine-grained, priority-aware admission controller, which manages the unified dynamic memory pool ($M_{dynamic}$). This scheduler acts as a crucial guardrail by enforcing a strict prioritization policy. It services the agent queue (Q_A) exhaustively, only considering a batch from the judge queue (Q_J) when the agent queue is empty or lacks sufficient memory for the next dispatch. This ensures that the deferrable, internal work of the judge can never block the critical path of a user-facing agent task. This two-level defense combines asymmetric partitioning for baseline efficiency with dynamic priority scheduling for worst-case protection, enabling robust, high-performance co-location.

5 Implementation and Discussion

Cortex is implemented atop vLLM [63], intercepting tool calls and integrating with agent applications (e.g., Search-R1 [32]). For co-location, we use CUDA MPS to let multiple processes share a single GPU. Even when local tools are low-latency, Cortex reduces cost by using semantic matching to avoid unnecessary tool calls and costs. We leave more advanced dynamic partitioning techniques (e.g., Green Contexts [49]) as future work.

Judge accuracy. Our framework bridges the gap between caching and semantic knowledge reuse by providing a unified cache abstraction for LLM agents, and it leverages powerful LLM-based techniques from prior research to ensure strong accuracy. The semantic judge can be easily fine-tuned or replaced for specific workloads, so its accuracy can be improved with minimal effort when needed. Given the rapid progress of small LLMs, we find this sufficient for practical use and view the judge as a pluggable component. In practice, Cortex maintains high precision via strong semantic models and periodic recalibration, with misses falling back to live fetches to preserve correctness.

6 Experiment

6.1 Experimental Setup

Testbed. We emulate a realistic cross-region deployment. The LLM agent and all Cortex components run on an on-premise H100 cluster; remote data services run in a separate region. For search agent, we use the public Google Cloud Search API whose per-request average latency ranges between 300–500ms depending on response length, consistent with prior work [40]. For coding, we use a self-deployed FAISS [17]-based RAG service [17] with 300 ms latency.

Models and workloads. We evaluate two agentic scenarios: AI-powered search and code generation. The search agent is Search-R1-7B [33, 34] (post-trained from Qwen-2.5 7B [2]); the coding agent is Qwen-3 8B [3]. Cortex uses Qwen-3 0.6B models for embedding and semantic judging [61, 62]. Cortex’s internal embedding and semantic judging functions are powered by the lightweight 0.6B models from the Qwen-3 family [61, 62].

Search workloads. Following the Google Trend access patterns in §2.3, we construct: (i) *Skewed workload:* extracting the Google-like head–tail access pattern in §2.3, we instantiate four search benchmarks—Zilliz-GPT [68], HotpotQA [7], Musique [8], and 2Wiki-Multi-Hop [6]. For each dataset, we apply k -means to the questions and keep 10 representative clusters, sampling ~ 250 questions and constructing the skewed popularity per dataset (1000 total) and (ii) *Trend-driven workload.* We capture 12-hour Google Trends time series for four topics, map them to HotpotQA questions, and compress them into a 10-minute trace to mimic sharp traffic spikes (§2.3).

SWE-Bench workload. We use a subset of widely-used SWE-

Bench_oracle [47] targeting the sqlfluff repo [11], where the agent resolves GitHub issues. Each request represents one issue. This stresses caching because multiple issues repeatedly access shared files, such as core files and project documents.

Baseline systems. To evaluate the end-to-end performance of an agent system, we compare three primary configurations, all built atop the vLLM framework [63] for fair comparison. We compare: *Agent_vanilla* (no cache; every request hits the remote API), *Agent_exact* (traditional exact-match KV cache), *Agent_Cortex* (our full system), and an ablation *Agent_ANN* that uses only ANN similarity. As *Agent_ANN* is impractical for production (§3.2), we restrict it to accuracy analysis (§6.6).

Metrics. We report throughput (req/s), latency (ms), cache hit rate (% served from cache), and operational cost (external API fees + GPU compute).

6.2 Overall Performance On Cache Ratio

Using the setup above, we vary cache size ratio and report throughput, hit rate, and latency across three workloads.

Skewed workload. We first evaluate Cortex on search benchmarks with head–tail popularity skew (§2.3) to show semantic matching benefits. As shown in Figure 7, Cortex outperforms the baselines across throughput, cache hit rate, and latency on all four datasets (Zilliz-GPT, HotpotQA, Musique, and 2Wiki). For instance, on Musique, *Agent_Cortex* achieves up to a $3.6\times$ higher throughput: *Agent_Cortex* sustains over 85% hit rates while the exact-match cache stays below 20%. This high hit rate reduces end-to-end latency by up to $4\times$ and minimizes reliance on external APIs, avoiding rate limits such as the Google Search API’s 100 queries per minute cap. Note that absolute latencies exceed raw network RTTs because requests to external APIs experience queueing and backoff under rate limits, which inflates latencies, as will be shown in §6.4. This performance gap stems from semantic diversity. *Agent_exact* fails on syntactic variations, yielding few cache hits, while *Agent_Cortex* groups semantically similar queries, converting semantic locality into high cache efficiency.

Trend-driven workload. We next evaluate *Agent_Cortex* on bursty workloads synthesized from Google Trends to test its adaptability. As shown in Figure 8, it achieves up to a $3.8\times$ throughput improvement over the *Agent_vanilla* baseline. This gain stems from its ability to maintain a high cache hit rate of nearly 95%, effectively absorbing the intense but temporary demand of each trending topic. This performance gain comes from the LCFU eviction policy, which is tailored for temporal dynamics. Traditional caches retain obsolete data from past trends, reducing effective capacity. In contrast, LCFU integrates an item’s staticity into its priority score. As a volatile topic’s popularity wanes, its cached entries are automatically deprioritized and become prime candidates for eviction. This mechanism ensures that cache space is continuously reclaimed for the next wave of content, which is crucial for maintaining high performance during bursty events.

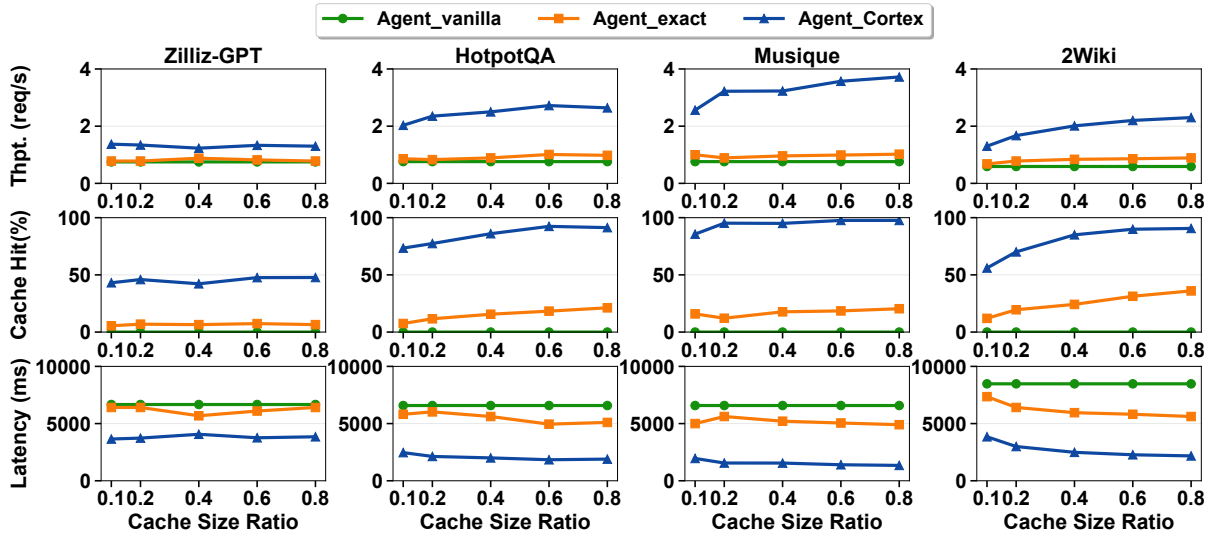


Figure 7: End-to-end agent serving throughput on skewed search workload under different cache ratio, zipfian-0.99 distribution




Figure 8: End-to-end agent serving throughput on trend-driven workload under different cache ratios




Figure 9: End-to-end agent serving throughput on SWE-Bench workload under different cache ratios

SWE-Bench workload. To demonstrate the generalizability of our approach, we evaluate Agent_Cortex on a code generation workload using the SWE-Bench. As shown in Figure 9, the results reveal a significant 20% throughput improvement over both Agent_vanilla and Agent_exact, driven by a cache hit rate approaching 45%. While more modest than the search results, these gains are highly impactful for the complex domain of software engineering and validate the versatility of our caching approach. The caching opportunity in this domain arises from shared file dependencies across tasks, such as when an agent resolves multiple GitHub issues within the




Figure 10: End-to-end Throughput Under Varying Request rate/concurrency on Musique Dataset

same repository. A traditional cache like Agent_exact is ineffective because it treats requests for different parts of the same file as distinct misses. Agent_Cortex, however, semantically identifies the shared file context, enabling it to serve these requests from the cache. Although the inherent diversity of coding tasks leads to a lower overall hit rate compared to search, serving nearly half of all file access requests locally provides a substantial reduction in latency and system load.

6.3 Scalability Under Concurrency

To assess the Cortex’s scalability under realistic workloads, we evaluate its end-to-end throughput with varying request concurrency. We use the Musique dataset at a fixed cache ratio of 0.4, a representative setting from our prior analysis. As shown in Figure 10, the results highlight Cortex’s significant scalability advantage, delivering up to a $5.7\times$ and $4.5\times$ throughput improvement comparing to the baselines. While the baseline systems quickly saturate, Agent_Cortex demonstrates strong scaling, achieving a throughput of 4.89 req/s at a request rate of 8. This represents a $4.5\times$ improvement over Agent_exact (1.09 req/s) and a $5.7\times$ improvement over Agent_vanilla (0.86 req/s) under heavy load. The throughput for both Agent_vanilla and Agent_exact plateaus early, at just




Figure 11: Per-request End-to-end Performance Breakdown




Figure 12: Data Retrieval Call and Retries Ratio

around 1 req/s. This is because their performance is fundamentally bound by the remote data retrieval. With low cache hit rates, nearly every request involves a high-latency external API call, waiting for these remote responses. Increasing concurrency merely leads to longer queues, not higher throughput. In contrast, Agent_Cortex’s throughput scales nearly linearly with concurrency up to a request rate of 8, where high cache efficiency serves most requests locally and fully utilizes GPU parallelism—driving the agent to hardware-level capacity. Beyond this point, Cortex maintains stable throughput (~5 req/s) even as the request rate increases to 32, with only minor degradation (<3%) at the highest load due to increased memory pressure. The result demonstrates the robustness of our co-location design which provides sufficient resources for both the agent and the semantic judge without degrading to vanilla behavior under stress.

6.4 Performance Breakdown

Latency analysis. To isolate pure request latency without rate-limit effects, we measure single-request breakdowns at low concurrency (Figure 11). While agent inference remains 0.6 s, Cortex cuts the total latency from 1.08 s to just 0.61s. This is achieved by almost entirely eliminating the 0.48 s external retrieval bottleneck that dominates the baseline’s execution time. In its place, Cortex introduces a minimal and predictable local overhead of only 0.05 s, comprising 0.02 s for cache retrieval and 0.03 s for judge validation. This trade-off validates our design principle: the marginal cost of intelligent local caching (0.05 s) is vastly outweighed by the savings from avoiding slow remote data fetches (0.48 s).

Cloud API rate limit and throughput analysis. We then study throughput under realistic load where rate limits matter (Figure 12). Beyond reducing single-request latency, Cortex significantly enhances system throughput and scalability by mitigating the critical bottleneck of external API rate lim-

Table 4: Normalized throughput comparison between w/o. API Rate Limit and w. API Rate Limit

	Without API Rate Limit	With API Rate Limit
Agent_vanilla	1	1
Agent_Cortex	1.5	4.16

Table 5: Cost and performance comparison across different configurations.

Metric	Agent_vanilla	Cortex w/o Sharing	Cortex
API Cost (\$)	6.5	6.5	0.64
GPU Cost (\$)	76	152	76
Total Cost (\$)	82.5	158.5	76.64
Thpt. (req/s)	0.87	4.74	4.89
Thpt./Cost (req/s/\$)	0.01	0.03	0.06

its. This is achieved by fundamentally reducing the system’s reliance on external services, as demonstrated in Figure 12. Due to high cache miss rates, the non-cached Agent_vanilla baseline generates approximately 1300 external API calls for the given task, leading to significant throttling and a high retry ratio of 25%. In stark contrast, Agent_Cortex slashes the API call count to just 103, a 92% reduction. This efficiency virtually cuts the retry ratio to a negligible 0.50%.

To quantify how this API traffic reduction translates into a throughput advantage, we run a controlled experiment as shown in Table 4. We use a self-deployed RAG data service with 300 ms here as we cannot cancel API rate limit for Google service. Even without rate limits, Cortex already provides a 1.5× throughput improvement over the vanilla system due to its inherent latency savings. However, when a realistic API rate limit is enforced, the advantage becomes far more pronounced, with Cortex achieving a 4.2× throughput gain over the throttled baseline. This comparison reveals that avoiding the rate-limiting bottleneck alone contributes an additional 2.8× improvement. By serving most requests locally, Cortex sidesteps rate-limit bottlenecks and scales robustly.

6.5 Cost Analysis

We now analyze cost-efficiency by evaluating three configurations under peak load on Musique (similar to §6.3): Agent_vanilla; Cortex w/o Sharing, which uses an extra GPU for the semantic judge; and the complete Cortex system. As detailed in Table 5, our analysis reveals that Cortex is not only faster but fundamentally more economical, delivering 6× more throughput per dollar than the vanilla baseline.

This cost advantage stems from resolving the API–compute trade-off. The Agent_vanilla system shows a total cost of \$82.5 to achieve a low throughput of 0.87 req/s. The Cortex w/o Sharing configuration illustrates the pitfall of a naive caching approach. While dramatically boosting throughput to 4.79 req/s, it does so by requiring a dedicated second GPU for the semantic judge, doubling the GPU cost from \$76 to \$152 and making it the most expensive setup (\$158.5). In contrast, our co-located design cuts API cost by >90% (to \$0.64) without extra hardware, keeping total spend at \$76.64




Figure 13: Generation quality comparison between Cortex and vanilla Search-R1 (Agent_vanilla)

Table 6: The benefits brought by LCFU.

Metric	Agent_LRU	Agent_LFU	LCFU
Cache hit (%)	0.88	0.89	0.86
Throughput (req/s)	2.14	2.16	2.35

Table 7: Co-location efficiency on H100.

Metric	Dedicated-2GPU	Co-located (MPS 80/20)
Throughput (req/s)	2.89	2.72
p99 latency (ms)	6601	7230

while retaining $\geq 95\%$ of the two-GPU throughput. Thus, Cortex delivers $\sim 5.6\times$ more performance at a similar cost, making it both fast and financially sustainable.

6.6 Deep Dive

Accuracy Analysis We measure correctness using Exact Match score [19, 32]. As shown in Figure 13, Cortex-w/o judge shows accuracy drops across datasets (e.g., StrategyQA: 0.69 vs. 0.79), demonstrating that naive ANN similarity can return wrong results. The full Cortex matches baseline accuracy—the Semantic Judge rejects non-equivalent candidates, confirming Cortex achieves caching efficiency without sacrificing correctness.

LCFU study. We tested our LCFU policy against LRU and LFU strategies on the HotpotQA workload, as shown in Table 6. While LFU achieved a higher hit rate (0.89 vs our 0.86), our policy delivered higher performance with up to 9% compared to LRU and LFU eviction policy, respectively. This trade-off is intentional: rather than maximizing hit rate, our policy prioritizes caching expensive-to-retrieve items, resulting in better system-wide latency reduction.

Co-location study. Co-locating the Agent and Semantic Judge on a single A100 via CUDA MPS (80/20), at a representative cache ratio (0.6), retains 94% of dedicated throughput (2.72 vs. 2.89 req/s) with a 9.5% increase in p99 latency, indicating a near-baseline efficiency.

Recalibration overhead. We also quantify the cost of threshold recalibration for the semantic judge using HotpotQA/Musique. Compared to a variant without recalibration, throughput decreases by only 2%. This bounded cost comes from brief offline validation on small samples to keep the

precision target, stabilizing accuracy under drift and preserving high-precision hits. In practice, the overhead is negligible relative to Cortex’s gains.

7 Other Related Work

Semantic prompt caching. Recent LLM serving systems [9, 24, 57, 65] like Google Apigee [24] match via embeddings to bypass inference; ContextCache [65] caches multi-turn responses; VectorQ [57] adapts thresholds for similar prompts. However, these approaches are largely internal to the inference process and typically for single-cloud/node, leaving the cross-region remote data bottleneck unaddressed. In contrast, Cortex operates at the external-data layer, caching knowledge retrieved via tool calls to cut cross-region latency and API costs. The two approaches are thus orthogonal and can be deployed together for complementary benefits.

Traditional and multi-cloud caching. Beyond semantic schemes, caches rely on LRU/LFU/TTL heuristics with Redis [55] or Memcached [45]. Cross-region deployments introduce latency, consistency, and cost trade-offs. Examples include Macaron [51] (auto-tuned tiers cutting data-lake costs under write-through), EVCache [46] (eventual consistency via Kafka), Azure Cache for Redis (passive cross-region replication with client co-location), and Alluxio [38] (global namespace for on-demand hot data). These architectures assume exact-match keys, while Cortex instead brings semantic-aware matching and adaptive policies to remote knowledge caching.

LLM-based agent memory. Finally, agent memory systems [29, 64, 67] improve LLM’s reasoning by retaining past context. For example, MemoryBank [67] stores conversational summaries to track entities; A-MEM [64] generalizes the memory structure to diverse tasks; Memory Sandbox [29] gives users explicit control. These systems give LLMs long-term memory to improve reasoning and choose the next action, but they don’t reduce the cost of executing tool calls. Cortex instead targets the external tool call itself, caching external results to make agent operation faster and cheaper.

8 Conclusion

We present Cortex, an agentic knowledge caching system designed to overcome the latency and cost challenges of LLM agents accessing external knowledge. Cortex’s novelty lies in its two-stage hybrid retrieval, combining ANN search with an LLM-powered Semantic Judge for semantic matching, and its resource-efficient co-location architecture. Our evaluation demonstrates that Cortex significantly improves performance and maintains accuracy, offering a scalable and cost-effective solution for LLM agent deployments.

Acknowledgments

We thank our shepherd Shay Vargaftik and anonymous reviewers for their valuable comments. This work is supported by the Singapore Ministry of Education, under Academic Research Fund Tier 1 grant T1 251RES2104.

References

- [1] Hyperbolic AI. Hyperbolic gpu marketplace: On-demand nvidia gpu rentals. <https://app.hyperbolic.ai/>, 2025. Accessed September 2025.
- [2] Alibaba Qwen Team. Qwen/qwen2.5-7b-instruct. <https://huggingface.co/Qwen/Qwen2.5-7B-Instruct>, 2025. Accessed: 2025-08-22.
- [3] Alibaba Qwen Team. Qwen/qwen3-8b. <https://huggingface.co/Qwen/Qwen3-8B>, 2025. Accessed: 2025-08-22.
- [4] Anthropic. Introducing the model context protocol. <https://www.anthropic.com/news/model-context-protocol>, 2025. Accessed: 2025-08-22.
- [5] Anthropic. Your code’s new collaborator. <https://www.anthropic.com/claude-code>, 2025. Accessed: 2025-07-22.
- [6] NLPiR Lab at RUC. 2wikimultihopqa: A multi-hop question answering dataset from wikipedia. https://huggingface.co/datasets/RUC-NLPiR/FlashRAG_datasets/tree/main/2wikimultihopqa, 2025. Part of FlashRAG datasets. Accessed: 2025-09-14.
- [7] NLPiR Lab at RUC. Hotpotqa (flashrag version): Multi-hop question answering dataset. https://huggingface.co/datasets/RUC-NLPiR/FlashRAG_datasets/tree/main/hotpotqa, 2025. Part of FlashRAG datasets. Accessed: 2025-09-14.
- [8] NLPiR Lab at RUC. Musique (flashrag version): Multi-hop question answering dataset. https://huggingface.co/datasets/RUC-NLPiR/FlashRAG_datasets/tree/main/musique, 2025. Part of FlashRAG datasets. Accessed: 2025-09-14.
- [9] Fu Bang. Gptcache: An open-source semantic cache for llm applications enabling faster answers and cost savings. In *Proceedings of the 3rd Workshop for Natural Language Processing Open Source Software (NLP-OSS 2023)*, pages 212–218, 2023.
- [10] Chroma Contributors. Chroma: Open-source search and retrieval for ai. <https://www.trychroma.com/>, 2025. Accessed September 2025.
- [11] SQLFluff Contributors. Sqlfluff: A modular sql linter and auto-formatter. <https://github.com/sqlfluff/sqlfluff>, 2025. Supports multiple SQL dialects and templated code including Jinja and dbt. Accessed: 2025-09-14.
- [12] Microsoft Corporation. The new bing: Ai-powered assistant for your search. <https://www.microsoft.com/en-us/edge/features/the-new-bing?form=MA13FJ>, 2025. Accessed September 2025.
- [13] Oracle Corporation. Mysql: The world’s most popular open source database. <https://www.mysql.com/>, 2025. Accessed September 2025.
- [14] CURSOR. The ai code editor. <https://cursor.com/agents>, 2025. Accessed: 2025-07-22.
- [15] DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucang Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanbiao Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xing kai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhanglei Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhenwen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. Deepseek-v3 technical report, 2025.
- [16] Ehsan Doostmohammadi, Tobias Norlund, Marco Kuhlmann, and Richard Johansson. Surface-based retrieval reduces perplexity of retrieval-augmented language models, 2023.
- [17] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. 2024.
- [18] Kimi Team et al. and Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, Yichen Feng, Kelin Fu, Bofei Gao, Hongcheng Gao, Peizhong Gao, Tong Gao, Xinran Gu, Longyu Guan, Haiqing Guo, Jianhang Guo, Hao Hu, Xiaoru Hao, Tianhong He, Weiran He, Wenyang He, Chao Hong, Yangyang Hu, Zhenxing Hu, Weixiao Huang, Zhiqi Huang, Zihao Huang, Tao Jiang, Zhejun Jiang, Xinyi Jin, Yongsheng Kang, Guokun Lai, Cheng Li, Fang Li, Haoyang Li, Ming Li, Wentao Li, Yanhao Li, Yiwei Li, Zhaowei Li, Zheming Li, Hongzhan Lin, Xiaohan

- Lin, Zongyu Lin, Chengyin Liu, Chenyu Liu, Hongzhang Liu, Jingyuan Liu, Junqi Liu, Liang Liu, Shaowei Liu, T. Y. Liu, Tianwei Liu, Weizhou Liu, Yangyang Liu, Yibo Liu, Yiping Liu, Yue Liu, Zhengying Liu, Enzhe Lu, Lijun Lu, Shengling Ma, Xinyu Ma, Yingwei Ma, Shaoguang Mao, Jie Mei, Xin Men, Yibo Miao, Siyuan Pan, Yebo Peng, Ruoyu Qin, Bowen Qu, Zeyu Shang, Lidong Shi, Shengyuan Shi, Feifan Song, Jianlin Su, Zhengyuan Su, Xinjie Sun, Flood Sung, Heyi Tang, Jiawen Tao, Qifeng Teng, Chensi Wang, Dinglu Wang, Feng Wang, Haiming Wang, Jianzhou Wang, Jiaying Wang, Jinhong Wang, Shengjie Wang, Shuyi Wang, Yao Wang, Yejie Wang, Yiqin Wang, Yuxin Wang, Yuzhi Wang, Zhaoji Wang, Zhengtao Wang, Zhexu Wang, Chu Wei, Qianqian Wei, Wenhao Wu, Xingzhe Wu, Yuxin Wu, Chenjun Xiao, Xiaotong Xie, Weimin Xiong, Boyu Xu, Jing Xu, Jinjing Xu, L. H. Xu, Lin Xu, Suting Xu, Weixin Xu, Xinran Xu, Yangchuan Xu, Ziyao Xu, Junjie Yan, Yuzi Yan, Xiaofei Yang, Ying Yang, Zhen Yang, Zhilin Yang, Zonghan Yang, Haotian Yao, Xingcheng Yao, Wenjie Ye, Zhuorui Ye, Bohong Yin, Longhui Yu, Enming Yuan, Hongbang Yuan, Mengjie Yuan, Haobing Zhan, Dehao Zhang, Hao Zhang, Wanlu Zhang, Xiaobin Zhang, Yangkun Zhang, Yizhi Zhang, Yongting Zhang, Yu Zhang, Yutao Zhang, Yutong Zhang, Zheng Zhang, Haotian Zhao, Yikai Zhao, Huabin Zheng, Shaojie Zheng, Jianren Zhou, Xinyu Zhou, Zaida Zhou, Zhen Zhu, Weiyu Zhuang, and Xinxing Zu. Kimi k2: Open agentic intelligence, 2025.
- [19] Hugging Face. Exact match score: Evaluation metric for text generation. https://huggingface.co/spaces/evaluate-metric/exact_match, 2025. Accessed September 2025.
- [20] Hao Fei, Yuan Yao, Zhuosheng Zhang, Fuxiao Liu, Ao Zhang, and Tat-Seng Chua. From multimodal llm to human-level ai: Modality, instruction, reasoning, efficiency and beyond. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024): Tutorial Summaries*, pages 1–8, 2024.
- [21] Ceph Foundation. Ceph: The future of storage. <https://ceph.io/en/>, 2025. Accessed September 2025.
- [22] In Gim, Guojun Chen, Seung-seob Lee, Nikhil Sarda, Anurag Khandelwal, and Lin Zhong. Prompt cache: Modular attention reuse for low-latency inference. *Proceedings of Machine Learning and Systems*, 6:325–338, 2024.
- [23] Google. Custom search api. <https://console.cloud.google.com/marketplace/product/google/customsearch.googleapis.com>, 2025. Accessed: 2025-08-22.
- [24] Google Cloud. Get started with semantic caching policies. <https://cloud.google.com/apigee/docs/api-platform/tutorials/using-semantic-caching-policies#:~:text=policies%20to%20enable%20intelligent%20response,latency%2C%20and%20lower%20operational%20costs>, 2025. Accessed: 2025-07-22.
- [25] gRPC Authors. Introduction to grpc. <https://grpc.io/docs/what-is-grpc/introduction/>, 2025. Accessed September 2025.
- [26] Vinit Kumar Gunjan, Monika Kumari, Amit Kumar, Allam Appa Rao, et al. Search engine optimization with google. *Asian Journal of Engineering and Applied Technology*, 1(1):36–42, 2012.
- [27] Luke Harsel. Google ai mode’s early adoption and seo impact. https://www.semrush.com/blog/google-ai-mode-seo-impact/?utm_source=chatgpt.com. Accessed: 2025-09-12.
- [28] Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, et al. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*, 3(4):6, 2023.
- [29] Ziheng Huang, Sebastian Gutierrez, Hemanth Kamana, and Stephen MacNeil. Memory sandbox: Transparent and interactive memory management for conversational agents, 2023.
- [30] Turbopuffer Inc. Turbopuffer: Serverless vector and full-text search built on object storage. <https://turbopuffer.com/>, 2025. Accessed September 2025.
- [31] Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnawamy, and Rohan Kadekodi. Diskann: Fast accurate billion-point nearest neighbor search on a single node. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [32] Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1: An efficient, scalable rl training framework for reasoning & search engine calling interleaved llms. <https://github.com/PeterGriffinJin/Search-R1>, 2025. Accessed: 2025-09-14.
- [33] Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. *arXiv preprint arXiv:2503.09516*, 2025.
- [34] Jin, Bowen and Zeng, Hansi and Yue, Zhenrui and Yoon, Jinsung and Arik, Sercan and Wang, Dong and Zamani, Hamed and Han, Jiawei. Searchr1-nq_hotpotqa_train-qwen2.5-7b-em-ppo. https://huggingface.co/PeterJinGo/SearchR1-nq_hotpotqa_train-qwen2.5-7b-em-ppo, 2025. Accessed: 2025-08-22.
- [35] Herve Jégou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128, 2011.
- [36] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626, 2023.
- [37] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474, 2020.

- [38] Haoyuan Li. *Alluxio: A virtual distributed file system*. University of California, Berkeley, 2018.
- [39] Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. Search-01: Agentic search-enhanced large reasoning models. *arXiv preprint arXiv:2501.05366*, 2025.
- [40] Chaofan Lin, Zhenhua Han, Chengruidong Zhang, Yuqing Yang, Fan Yang, Chen Chen, and Lili Qiu. Parrot: Efficient serving of {LLM-based} applications with semantic variable. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pages 929–945, 2024.
- [41] Yuhan Liu, Hanchen Li, Yihua Cheng, Siddhant Ray, Yuyang Huang, Qizheng Zhang, Kuntai Du, Jiayi Yao, Shan Lu, Ganesh Ananthanarayanan, et al. Cachegen: Kv cache compression and streaming for fast large language model serving. In *Proceedings of the ACM SIGCOMM 2024 Conference*, pages 38–56, 2024.
- [42] Google LLC. Google ai mode: A new way to search, whatever’s on your mind. <https://search.google/ways-to-search/ai-mode/>, 2025. Accessed September 2025.
- [43] Google LLC. Google programmable search engine: Custom search powered by google. <https://programmablesearchengine.google.com/about/>, 2025. Accessed September 2025.
- [44] Peixian Ma, Xialie Zhuang, Chengjin Xu, Xuhui Jiang, Ran Chen, and Jian Guo. Sql-r1: Training natural language to sql reasoning model by reinforcement learning. *arXiv preprint arXiv:2504.08600*, 2025.
- [45] Memcached. Memcached. <https://memcached.org/>, 2025. Accessed: 2025-07-22.
- [46] Netflix. Evcache. <https://github.com/Netflix/EVCache>, 2025. Accessed: 2025-07-22.
- [47] Princeton NLP. Swe-bench oracle dataset. https://huggingface.co/datasets/princeton-nlp/SWE-bench_oracle, 2023. Accessed: 2025-09-14.
- [48] NVIDIA. Multi-process service. <https://docs.nvidia.com/deploy/mps/index.html>, 2025. Accessed: 2025-07-22.
- [49] NVIDIA Corporation. CUDA Driver API: Green Contexts. https://docs.nvidia.com/cuda/cuda-driver-api/group__CUDA__GREEN__CONTEXTS.html, 2025. Accessed: 2025-12-31.
- [50] OpenAI. Api pricing. <https://openai.com/api/pricing/>, 2025. Accessed: 2025-09-07.
- [51] Hojin Park, Ziyue Qiu, Gregory R Ganger, and George Amvrosiadis. Reducing cross-cloud/region costs with the auto-configuring macaron cache. In *Proceedings of the ACM SIGOPS 30th Symposium on Operating Systems Principles*, pages 347–368, 2024.
- [52] Shishir G. Patil, Tianjun Zhang, Xin Wang, and Joseph E. Gonzalez. Gorilla: Large language model connected with massive apis. *arXiv preprint arXiv:2305.15334*, 2023.
- [53] Sarah Perez. Google’s ai overviews have 2b monthly users, ai mode 100m in the us and india. https://techcrunch.com/2025/07/23/googles-ai-overviews-have-2b-monthly-users-ai-mode-100m-in-the-us-and-india/?utm_source=chatgpt.com. Accessed: 2025-09-12.
- [54] Partha Pratim Ray. A survey on model context protocol: Architecture, state-of-the-art, challenges and future directions. *Authorea Preprints*, 2025.
- [55] Redis. Redis. <https://redis.io/>, 2025. Accessed: 2025-07-22.
- [56] Sachin Agarwal. Public cloud inter-region network latency as heat-maps. <https://medium.com/@sachinkagarwal/public-cloud-inter-region-network-latency-as-heat-maps-134e22a5ff19>, 2025. Accessed: 2025-07-22.
- [57] Luis Gaspar Schroeder, Shu Liu, Alejandro Cuadron, Mark Zhao, Stephan Krusche, Alfons Kemper, Matei Zaharia, and Joseph E Gonzalez. Adaptive semantic prompt caching with vectorq. *CoRR*, 2025.
- [58] LLC SerpApi. Serpapi: Real-time search engine results api. <https://serpapi.com/>, 2025. Accessed September 2025.
- [59] Huatong Song, Jinhao Jiang, Yingqian Min, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, and Ji-Rong Wen. R1-searcher: Incentivizing the search capability in llms via reinforcement learning. *arXiv preprint arXiv:2503.05592*, 2025.
- [60] Facebook Database Engineering Team. Rocksdb: A persistent key-value store for fast storage. <https://github.com/facebook/rocksdb>, 2025. Accessed September 2025.
- [61] Qwen Team. Qwen3-embedding-0.6b: A multilingual text embedding model. <https://huggingface.co/Qwen/Qwen3-Embedding-0.6B>, 2025. Supports over 100 languages and customizable embedding dimensions. Accessed: 2025-09-14.
- [62] Qwen Team. Qwen3-reranker-0.6b: A multilingual text reranking model. <https://huggingface.co/Qwen/Qwen3-Reranker-0.6B>, 2025. Supports over 100 languages and customizable instructions for retrieval tasks. Accessed: 2025-09-14.
- [63] vLLM Project Contributors. vllm: A high-throughput and memory-efficient inference and serving engine for llms. <https://github.com/vllm-project/vllm>, 2025. Originally developed at Sky Computing Lab, UC Berkeley. Accessed: 2025-09-14.
- [64] Wujiang Xu, Kai Mei, Hang Gao, Juntao Tan, Zujie Liang, and Yongfeng Zhang. A-mem: Agentic memory for llm agents, 2025.
- [65] Jianxin Yan, Wangze Ni, Lei Chen, Xuemin Lin, Peng Cheng, Zhan Qin, and Kui Ren. Contextcache: Context-aware semantic cache for multi-turn queries in large language models. *arXiv preprint arXiv:2506.22791*, 2025.
- [66] Danna Zheng, Mirella Lapata, and Jeff Z Pan. Large language models as reliable knowledge bases? 2024.
- [67] Wanjun Zhong, Lianghong Guo, Qiqi Gao, He Ye, and Yanlin Wang. Memorybank: Enhancing large language models with long-term memory, 2023.
- [68] zilliztech. Zilliz gpt cache datasets. https://github.com/zilliztech/GPTCache/blob/main/examples/benchmark/similiar_qgp_full.json.gz, 2025. Accessed: 2025-08-22.