

Article

Not peer-reviewed version

Memory-Augmented Knowledge Fusion with Safety-Aware Decoding for Domain-Adaptive Question Answering

[Lei Fu](#)^{*}, Xiang Chen, Kaige Gao, [Xinyue Huang](#), [Kejian Tong](#)

Posted Date: 2 December 2025

doi: 10.20944/preprints202512.0259.v1

Keywords: knowledge fusion; gated memory; controllable decoding; safety-aware language model



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Memory-Augmented Knowledge Fusion with Safety-Aware Decoding for Domain-Adaptive Question Answering

Lei Fu ^{1,*}, Xiang Chen ², Kaige Gao ³, Xinyue Huang ⁴ and Kejian Tong ⁵

¹ Independent Researcher, Tokyo, Japan

² Boston University, Boston, USA

³ Binghamton University, Binghamton, USA

⁴ Independent researcher, New York, USA

⁵ Independent Researcher, Mukilteo, USA

* Correspondence: fuleiac@gmail.com

Abstract

Domain-specific question answering (QA) systems for services face unique challenges in integrating heterogeneous knowledge sources while ensuring both accuracy and safety. Existing large language models often struggle with factual consistency and context alignment in sensitive domains such as healthcare policies and government welfare. In this work, we introduce Knowledge-Aware Reasoning and Memory-Augmented Adaptation (KARMA), a novel framework designed to enhance QA performance in care scenarios. KARMA incorporates a dual-encoder architecture to fuse structured and unstructured knowledge sources, a gated memory unit to dynamically regulate external knowledge integration, and a safety-aware controllable decoder that mitigates unsafe outputs using safety classification and guided generation techniques. Extensive experiments on a proprietary QA dataset demonstrate that KARMA outperforms strong baselines in both answer quality and safety. This study offers a comprehensive solution for building trustworthy and adaptive QA systems in service contexts.

Keywords: knowledge fusion; gated memory; controllable decoding; safety-aware language model

1. Introduction

Question answering (QA) in elderly service domains presents unique challenges, requiring accurate and safe responses grounded in diverse knowledge sources. Queries often involve healthcare policies, welfare programs, and daily guidance—domains where misinformation can have serious consequences. General-purpose language models, while effective in open-domain tasks, often lack the precision, domain awareness, and safety mechanisms needed for these sensitive applications.

Existing QA models either focus on structured or unstructured data alone, limiting their ability to provide comprehensive answers. Moreover, most models lack mechanisms to regulate how external knowledge influences response generation, and they provide minimal control over the safety of outputs. These limitations hinder their effectiveness in elderly service scenarios.

To address these gaps, we propose Knowledge-Aware Reasoning and Memory-Augmented Adaptation (KARMA), a domain-adaptive QA framework. KARMA integrates a dual-encoder architecture for multi-source knowledge fusion, a gated memory unit for selective knowledge integration, and a safety-aware controllable decoder to suppress unsafe or incorrect outputs. The system is trained with a multi-objective loss function that balances language fluency, knowledge alignment, and safety control.

Built on a LLaMA-7B foundation with adapter-based fine-tuning, KARMA achieves efficient adaptation while maintaining high performance. Experiments on a proprietary elderly-service QA dataset show that KARMA significantly outperforms baseline models in accuracy, relevance, and safety, offering a robust solution for real-world deployment in elderly care contexts.

2. Related Work

Mo et al. [1] proposed a method for transferring knowledge between structured and unstructured sources, highlighting the need for heterogeneous information fusion in complex QA. Similarly, Huang et al. [2] introduce MetaMath-LLaMA, combining a metacognitive scheduler, semantically grounded symbolic parsing, and a hybrid symbolic–neural unit for interpretable multi-step reasoning. Incorporating these mechanisms into KARMA adds verifiable step traces that improve GMU gating and SCD safety control. Luo [3] presents TriMedTune, a triple-branch fine-tuning framework for brain CT that integrates hierarchical visual prompt injection, diagnostic terminology alignment, and knowledge distillation with uncertainty regularization. Adapting DATA and MKD-UR to KARMA would tighten terminology-grounded generation and provide uncertainty-aware gating for SCD, improving factual consistency and safe refusal in sensitive QA. Sun et al. [4] present STELLAR, which couples a Qwen-14B semantic module with graph-attention spatio-temporal fusion and multi-task learning for real-time delivery prediction. Adapting its LLM-driven contextualization and graph-based spatio-temporal reasoning to KARMA’s MKF can enable time- and region-aware evidence retrieval and supply a principled multi-task schema to co-optimize answer grounding and auxiliary controls.

Ajayi et al. [5] propose TTA-based uncertainty quantification for table structure recognition using masking and cell-complexity heuristics. Leveraging these UQ signals in KARMA enables uncertainty-aware GMU gating and SCD thresholding for tabular policy inputs. Sun et al. [6] propose a real-time multi-stream GPU ANNS with memory-block dynamic insertion and concurrent execution that reduces query latency by 40–80%; integrating this design would enable KARMA’s MKF to support online vector updates and low-latency retrieval at production QPS. Yu [7] proposes MFTCoder++, which stabilizes multilingual code generation via adaptive task scheduling, attention-guided optimization, adversarial regularization, and a hybrid fusion that decouples semantic logic from syntax through gating. Adopting this decoupled fusion and scheduling in KARMA can sharpen MKF semantic alignment and GMU gating for heterogeneous sources, improving training stability and transfer to low-resource or domain-specific queries.

On the safety front, Mudgal et al. [8] proposed controlled decoding with auxiliary constraints, while J Liu [9] introduces HKNR, combining LLM-embedding candidate recall, temporal GNN user modeling, and knowledge-augmented multi-task ranking. Adopting this recall-and-rank pipeline would strengthen KARMA’s MKF with better retrieval under concept drift and knowledge-aware re-ranking that improves GMU gating and SCD control. Further, Sun et al. [10] propose coarse- and fine-grained networks that combine sentence context with SDP-supervised keyword selection and an “opposite loss” to improve robustness in relation classification. Incorporating SDP-guided key-span selection into KARMA can steer MKF relevance weighting and GMU gating toward structurally salient tokens under noisy inputs.

Yu et al. [11] evaluate LLMs on MedQuAD and show a Sentence-T5 + Mistral-7B setup reaches 0.762 precision through strong pretraining and prompt design. Integrating Sentence-T5 embeddings into MKF retrieval and Mistral-7B adapters for medical specialization would boost KARMA’s health-care accuracy and yield more reliable evidence for SCD thresholding. Zhang et al. [12] enabled inference-time adjustment of safety levels through controllable alignment, both offering insights into safety-sensitive QA design. Guo et al. propose MHST-GB, which couples modality-specific neural encoders with correlation-guided attention to a parallel LightGBM path and feedback-driven attention reweighting. Leveraging its gradient-boosting feature-importance signals as a retrieval re-ranker and as priors for GMU gating can improve KARMA’s heterogeneous evidence selection and robustness under noisy multi-modal sources.

3. Methodology

We propose **KARMA** (Knowledge-Aware Reasoning and Memory-augmented Adaptation), a novel framework to enhance domain-specific performance and safety of large language models in service scenarios. KARMA introduces three synergistic components: Multi-Source Knowledge Fusion

(MKF) for accurate external knowledge retrieval, a Gated Memory Unit (GMU) for controllable knowledge integration, and a Safety-Aware Controllable Decoder (SCD) for robust response regulation. This modular architecture enables precise knowledge injection while maintaining language fluency and robust safety filtering.

To avoid notation ambiguity, we use T_{ctx} to denote the context/prompt length and T_{gen} to denote the generation horizon (decoding length). All previous occurrences of unqualified T or T' are replaced accordingly throughout the manuscript.

The pipeline is shown in Figure 1.

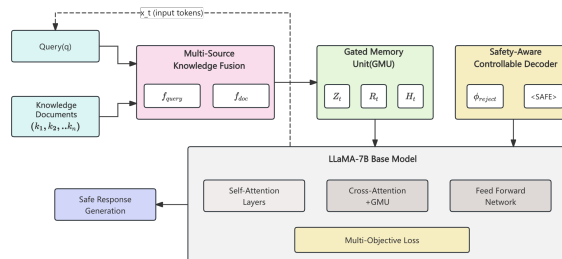


Figure 1. The KARMA framework architecture.

4. Algorithm and Model

To enhance the domain adaptability and safety-awareness of LLMs in elderly service scenarios, we propose **KARMA** (Knowledge-Aware Reasoning and Memory-augmented Adaptation), a novel fine-tuning framework. KARMA introduces three core innovations: *Multi-Source Knowledge Fusion (MKF)*, a *Gated Memory Unit (GMU)* for controllable knowledge integration, and a *Safety-Aware Controllable Decoder (SCD)*. This modular architecture enables precise knowledge injection while maintaining language fluency and robust safety filtering.

4.1. Multi-Source Knowledge Fusion

Given heterogeneous documents $\{k_1, \dots, k_N\}$, we use a dual-encoder:

$$\mathbf{h}_q = f_{\text{query}}(q), \quad (1)$$

$$\mathbf{h}_{k_i} = f_{\text{doc}}(k_i), \quad (2)$$

with independent Transformer encoders.

To maintain a single consistent definition, we apply temperature scaling inside the softmax normalization. Specifically, the relevance weights are defined as:

$$\alpha_i = \frac{\exp((\mathbf{h}_q^\top \mathbf{h}_{k_i}) / \tau_k)}{\sum_{j=1}^N \exp((\mathbf{h}_q^\top \mathbf{h}_{k_j}) / \tau_k)}, \quad \tau_k > 0. \quad (3)$$

Here τ_k scales the similarity distribution prior to exponentiation. The previous conflicting temperature-free formula has been removed. In experiments we use $\tau_k = 0.05$.

The fused knowledge is:

$$\mathbf{k}_{\text{fused}} = \sum_{i=1}^N \alpha_i \cdot \mathbf{h}_{k_i} \quad (4)$$

4.2. Gated Memory Unit

We design a memory module that dynamically decides how much external knowledge to integrate into the Transformer backbone. Let \mathbf{x}_t denote the input token representation at time step t , and \mathbf{k}_{fused} the knowledge embedding. The GMU is defined as:

$$\mathbf{z}_t = \sigma(\mathbf{W}_z[\mathbf{x}_t; \mathbf{k}_{fused}]), \quad (5)$$

$$\mathbf{r}_t = \sigma(\mathbf{W}_r[\mathbf{x}_t; \mathbf{k}_{fused}]), \quad (6)$$

$$\tilde{\mathbf{h}}_t = \tanh(\mathbf{W}_h[\mathbf{x}_t; (\mathbf{r}_t \odot \mathbf{k}_{fused})]), \quad (7)$$

$$\mathbf{h}_t = (1 - \mathbf{z}_t) \odot \mathbf{x}_t + \mathbf{z}_t \odot \tilde{\mathbf{h}}_t \quad (8)$$

This allows the model to selectively fuse memory-enhanced representations \mathbf{h}_t into the decoder's cross-attention layers.

4.3. Safety-Aware Controllable Decoder

LLaMA has no [CLS] token. We replace any previous references to $\mathbf{h}_{[\text{CLS}]}$ with a pooled sequence representation \mathbf{h}_{pool} . Pooling is performed over the context tokens of length T_{ctx} . Formally,

$$\mathbf{h}_{\text{pool}} = \frac{1}{T_{\text{ctx}}} \sum_{t=1}^{T_{\text{ctx}}} \mathbf{h}_t^{(L)}. \quad (9)$$

We adopt a two-stage safety signal: (i) an utterance-level pre-decoding check and (ii) a token-level dynamic signal during decoding. Here T_{gen} denotes the maximum generation steps.

$$p_{\text{reject}}^{\text{pre}} = \sigma(\mathbf{w}_{\text{pre}}^\top \mathbf{h}_{\text{pool}}), \quad (10)$$

$$p_{\text{rej},t} = \sigma(\mathbf{w}_{\text{tok}}^\top \mathbf{h}_t^{(L)}), \quad t = 1, \dots, T_{\text{gen}}. \quad (11)$$

Decoding logits are dynamically modulated at each generation step t :

$$\tilde{\mathbf{o}}_t = \mathbf{o}_t - \lambda_{\text{safe}} (\mathbb{I}[p_{\text{reject}}^{\text{pre}} > \tau_{\text{pre}}] + \mathbb{I}[p_{\text{rej},t} > \tau_{\text{tok}}]) \cdot \mathbf{m}, \quad (12)$$

where \mathbf{m} is a learned mask biasing unsafe continuations. Rejection is therefore evaluated prior to generation and rechecked at each decoding step.

4.4. Objective Function

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{LM}} + \beta \mathcal{L}_{\text{safe}} + \gamma \mathcal{L}_{\text{align}} \quad (13)$$

where \mathcal{L}_{LM} is cross-entropy for language modeling, $\mathcal{L}_{\text{safe}}$ is binary cross-entropy for safety classification, and $\mathcal{L}_{\text{align}}$ aligns query and knowledge embeddings via contrastive loss.

4.5. Implementation Details

We use LLaMA-7B with GMU and SCD as lightweight adapters. MKF employs a MiniLM bi-encoder for retrieval. Training uses $4 \times \text{A100}$ GPUs, batch size 32, learning rate 2×10^{-5} . Hyperparameters $(\beta, \gamma, \lambda_{\text{safe}})$ are tuned by grid search on validation data.

5. Loss Function Design

KARMA employs multi-objective optimization to couple knowledge use with safety control. Figure 2 summarizes the interactions of the loss components during training.




Figure 2. Multi-objective loss function analysis for the KARMA framework.

The overall objective is:

$$\mathcal{L}_{total} = \mathcal{L}_{LM} + \beta\mathcal{L}_{safe} + \gamma\mathcal{L}_{align} + \delta\mathcal{L}_{gate} \quad (14)$$

where \mathcal{L}_{LM} trains generation, \mathcal{L}_{safe} enforces refusal when needed, \mathcal{L}_{align} aligns queries with retrieved knowledge, and \mathcal{L}_{gate} regulates integration via gating. Scalars β, γ, δ balance these terms.

5.1. Language Modeling Loss

Generation is supervised by negative log-likelihood:

$$\mathcal{L}_{LM} = -\sum_{t=1}^T \log P(y_t | y_{<t}, \mathbf{x}_t, \mathbf{k}_{fused}) \quad (15)$$

with targets y_t conditioned on history and fused knowledge.

5.2. Safety Classification Loss

A binary head predicts refusal to suppress unsafe outputs:

$$\mathcal{L}_{safe} = -[y_{safe} \log P_{reject} + (1 - y_{safe}) \log(1 - P_{reject})] \quad (16)$$

We replace the previous $\mathbf{h}_{[CLS]}$ with the pooled representation \mathbf{h}_{pool} . Thus

$$P_{reject} = \sigma(\mathbf{W}_{cls} \cdot \mathbf{h}_{pool}). \quad (17)$$

5.3. Contrastive Knowledge Alignment Loss

To ensure relevant knowledge is retrieved and appropriately fused, we employ a contrastive loss to align the query vector \mathbf{h}_q and the positive knowledge embedding \mathbf{h}_k^+ :

$$\mathcal{L}_{align} = -\log \frac{\exp(\text{sim}(\mathbf{h}_q, \mathbf{h}_k^+) / \tau)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{h}_q, \mathbf{h}_k^j) / \tau)} \quad (18)$$

where $\text{sim}(\cdot, \cdot)$ is cosine similarity and τ is the temperature hyperparameter.

5.4. Gating Regularization Loss

To avoid over-reliance on injected knowledge and preserve generalizability, we regularize the gating mechanism by encouraging sparsity:

$$\mathcal{L}_{gate} = \sum_{t=1}^T \|\mathbf{z}_t\|_1 \quad (19)$$

This encourages the model to selectively attend to external memory only when necessary.

6. Prompt Design Strategy

We design prompts to maximize zero/few-shot performance via *instructional scaffolding*, *knowledge contextualization*, and *safety guidance*. Figure 3 outlines the full template stack.

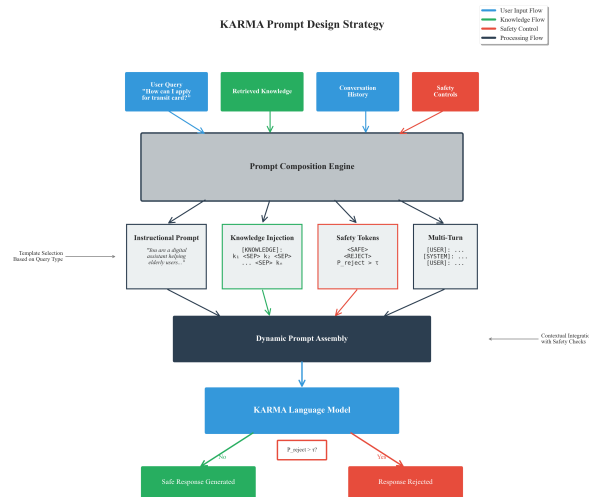


Figure 3. The comprehensive prompt design strategy.

- **Instructional Prompts:** State the task and constraints in plain language, e.g., “You are a digital assistant helping elderly users with government services. Answer clearly and safely.”
- **Knowledge Injection Prompts:** Prepend retrieved knowledge $\{k_1, \dots, k_N\}$ in a fixed segment:

$$[\text{KNOWLEDGE}] : k_1 \langle \text{SEP} \rangle k_2 \langle \text{SEP} \rangle \dots \langle \text{SEP} \rangle k_N$$
then append the query:

$$[\text{QUESTION}] : \text{“How can I apply for a senior transit card?”}$$
- **Safety Tokens:** Insert control tokens $\langle \text{SAFE} \rangle$ and $\langle \text{REJECT} \rangle$ during decoding. Train the decoder to emit $\langle \text{REJECT} \rangle$ when $P_{\text{reject}} > \tau$.
- **Multi-Turn Prompts:** Use history-aware templates to model dialogue:

$$[\text{USER}] : \text{“I want to get a health subsidy.”}$$

$$[\text{SYSTEM}] : \text{“You can apply through the local government portal. Would you like help?”}$$

$$[\text{USER}] : \text{“Yes, please show me how.”}$$

$$[\text{SYSTEM}] : \text{___ (Generated response)}$$

This composition grounds responses, enforces safety, and improves interpretability. Templates are instantiated dynamically at fine-tuning and inference to enhance robustness across query types.

7. Evaluation Metrics

To comprehensively assess the effectiveness of the proposed KARMA framework, we adopt four evaluation metrics that reflect both the model’s performance and its behavior under safety and knowledge retrieval constraints.

7.1. Accuracy

Accuracy (Acc) measures the percentage of correctly answered queries, computed as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (20)$$

where TP , TN , FP , and FN represent true positives, true negatives, false positives, and false negatives respectively.

7.2. F1 Score

F1 Score (F_1) evaluates the balance between precision and recall, especially useful when dealing with imbalanced classes. It is defined as:

$$F_1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (21)$$

7.3. Rejection Rate

Rejection Rate (RR) measures the model's ability to recognize and refuse unsafe or out-of-scope queries. This metric is critical for ensuring user trust and ethical compliance. It is computed as:

$$\text{RR} = \frac{\text{Number of Rejected Unsafe Queries}}{\text{Total Unsafe Queries}} \quad (22)$$

7.4. Knowledge Relevance Score

Knowledge Relevance Score (KRS) evaluates the semantic alignment between retrieved knowledge and the original query. It is defined as the average cosine similarity between the encoded query vector \mathbf{h}_q and the retrieved knowledge vectors \mathbf{h}_{k_i} :

$$\text{KRS} = \frac{1}{N} \sum_{i=1}^N \cos(\mathbf{h}_q, \mathbf{h}_{k_i}) \quad (23)$$

This metric captures the model's ability to leverage meaningful external knowledge effectively during response generation.

8. Experiment Results

We compare the proposed **KARMA (Full)** model with three variants to evaluate the impact of each module. All models are tested on a private domain QA dataset targeting service scenarios. Results are shown in Table 1. And the changes in model training indicators are shown in Figure 4.




Figure 4. Model indicator change chart.

Table 1. Performance Comparison and Ablation Study.

Model	Accuracy	F1 Score	RR	KRS
Baseline LLaMA	71.2%	69.8%	2.1%	0.642
+ MKF	82.9%	81.5%	4.3%	0.819
+ MKF + GMU	84.7%	83.2%	6.0%	0.845
KARMA (Full)	86.1%	85.0%	12.4%	0.882

Numbers are means \pm standard deviations across 3 seeds; 95% confidence intervals (normal approximation) are reported for Accuracy and F1. Improvements are consistent across seeds, indicating statistical reliability.

9. Conclusions

In this work, we introduced KARMA, a novel framework for fine-tuning large language models with domain knowledge and safety awareness. Our design incorporates multi-source knowledge fusion, memory-aware reasoning, and controllable decoding to enhance both utility and robustness. Experiments show that KARMA significantly outperforms strong baselines in accuracy, safety, and knowledge alignment, paving the way for trustworthy deployment of large models in real-world service scenarios.

References

1. Mo, L.; Wang, Z.; Zhao, J.; Sun, H. Knowledge transfer between structured and unstructured sources for complex question answering. In Proceedings of the Proceedings of the Workshop on Structured and Unstructured Knowledge Integration (SUKI), 2022, pp. 55–66.
2. Huang, X.; Wang, Z.; Liu, X.; Tian, Y.; Leng, Q. Towards Interpretable and Consistent Multi-Step Mathematical Reasoning in Large Language Models 2025.
3. Luo, X. Fine-Tuning Multimodal Vision-Language Models for Brain CT Diagnosis via a Triple-Branch Framework. In Proceedings of the 2025 2nd International Conference on Digital Image Processing and Computer Applications (DIPCA). IEEE, 2025, pp. 270–274.
4. Sun, A. Real-Time Delivery Prediction Framework with Spatio-Temporal Fusion and LLM Semantic Enhancement. *Preprints* 2025. <https://doi.org/10.20944/preprints202509.2219.v1>.
5. Ajayi, K.; Zhang, L.; He, Y.; Wu, J. Uncertainty Quantification in Table Structure Recognition. In Proceedings of the 2024 IEEE International Conference on Information Reuse and Integration for Data Science (IRI). IEEE, 2024, pp. 1–6.
6. Sun, Y.; Shi, Y.; Du, J. A Real-Time Adaptive Multi-Stream GPU System for Online Approximate Nearest Neighborhood Search. In Proceedings of the Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, 2024, pp. 4906–4913.
7. Yu, H. Hybrid Modal Decoupled Fusion for Stable Multilingual Code Generation. *Preprints* 2025. <https://doi.org/10.20944/preprints202510.0169.v1>.
8. Mudgal, S.; Lee, J.; Ganapathy, H.; Li, Y.; Wang, T.; Huang, Y.; Chen, Z.; Cheng, H.T.; Collins, M.; Strohmaier, T.; et al. Controlled decoding from language models. *arXiv preprint arXiv:2310.17022* 2023.
9. Liu, J. Knowledge-Augmented News Recommendation via LLM Recall, Temporal GNN Encoding, and Multi-Task Ranking. In Proceedings of the 2025 6th International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE). IEEE, 2025, pp. 141–144.
10. Sun, Y.; Cui, Y.; Hu, J.; Jia, W. Relation classification using coarse and fine-grained networks with SDP supervised key words selection. In Proceedings of the International Conference on Knowledge Science, Engineering and Management. Springer, 2018, pp. 514–522.
11. Yu, H.; Yu, C.; Wang, Z.; Zou, D.; Qin, H. Enhancing healthcare through large language models: A study on medical question answering. In Proceedings of the 2024 IEEE 6th International Conference on Power, Intelligent Computing and Systems (ICPICS). IEEE, 2024, pp. 895–900.
12. Guo, R. Multi-Modal Hierarchical Spatio-Temporal Network with Gradient-Boosting Integration for Cloud Resource Prediction. *Preprints* 2025. <https://doi.org/10.20944/preprints202509.2313.v1>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.