

Toward a Unified Multi-source Retrieval-Augmented Generation: Integrating RAG-VR and iRAG-MM Approaches

Faisal Alanqoodi
Middle East College (MEC)
faisalalanqood@gmail.com

Abstract—Retrieval-Augmented Generation (RAG) has emerged as a promising paradigm to enhance Large Language Models (LLMs) by coupling them with external knowledge sources. Two lines of recent research, RAG-VR (a method focusing on structured versioning mechanisms in end-to-end pipelines) and iRAG-MM (a multi-domain, multi-source retrieval approach with intelligent orchestration), provide complementary solutions to address the challenges of dynamic data dependencies, multi-source interactions, and the necessity for versioned knowledge integration. In this paper, we unify the essential ideas from RAG-VR and iRAG-MM into a single comprehensive framework. Specifically, we propose a system that (1) leverages versioning of intermediate datasets for robust pipeline coordination; (2) orchestrates multi-source retrieval with domain-specific adapters; and (3) employs a confidence-based fusion of evidence, thus ensuring consistency, interpretability, and adaptiveness across a spectrum of knowledge-intensive tasks. Through theoretical arguments, design considerations, and experimental insights, we illustrate how the unified system addresses real-world demands for personalization, factual accuracy, incremental updating of knowledge, and efficient co-development across teams. This integrated perspective not only highlights the synergy between version-aware RAG pipelines and multi-domain orchestrations but also paves the way for future advancements in retrieval-augmented generative models, bridging the gap between production environments and experimental flexibility.

I. INTRODUCTION

Large Language Models (LLMs) have demonstrated remarkable performance in various tasks ranging from open-domain question answering to more specialized domains such as legal or medical dialogues. Despite their strengths, they face persistent issues of hallucination, difficulty in keeping up with current knowledge, and a limited ability to adapt seamlessly to user-specific or domain-specific contexts. To address these challenges, **Retrieval-Augmented Generation (RAG)** fuses LLMs with external data repositories, enabling the model to fetch up-to-date or high-fidelity evidence before generating responses. The resulting pipeline, however, can be quite complex, involving multiple intermediate processing steps, multiple knowledge sources, and frequent changes in code or hyperparameters.

Two frameworks, **RAG-VR** and **iRAG-MM**, have separately tackled different aspects of this challenge:

- 1) **RAG-VR**: Emphasizes a *versioning* perspective on intermediate datasets, thus enabling multiple teams and

systems to concurrently run or update machine learning pipelines without synchronization overhead. This approach focuses on carefully managing data “derivations” and ensuring reusability of intermediate stages.

- 2) **iRAG-MM**: Concentrates on orchestrating *multiple sources of knowledge* with a specialized “intelligent orchestration” step, domain-specific adapters, and a confidence-based fusion strategy. This approach addresses multi-domain tasks that require more than a single knowledge source, ensuring that LLMs retrieve the right evidence from the right source and that collisions or redundancies are minimized.

A. Objective and Scope

This research attempts to merge the main ideas from both frameworks into a single, integrated system, called **UniRAG-VR-MM** (working name). By combining advanced versioning strategies for pipeline stages with dynamic multi-source retrieval orchestration, the proposed system aims to:

- Provide a robust solution for pipeline-level concurrency, partial re-computation, and parallel deployment.
- Offer a flexible approach for retrieving from multiple sources on demand (with domain adapters) and merging these outputs in a confidence-driven manner.
- Achieve synergy between “version-awareness” (handling code/data evolutions gracefully) and “multi-domain orchestration” (improving coverage and specificity for queries spanning different knowledge bases).

B. Contributions

- 1) **Unified Pipeline Model**: Presenting a pipeline design that adopts version-based dataset derivation (RAG-VR style) and extends it with multi-domain orchestrations (iRAG-MM style).
- 2) **Mathematical Formulation of Sources**: Proposing a formula to quantify confidence-based weighting of evidence, bridging the versioning constraints.
- 3) **Implementation and Practical Aspects**: Detailing how the system can be deployed in real-world setups. This includes storing partial derivations, generating source call orders, and updating knowledge incrementally.
- 4) **Extensive Analysis**: Demonstrating through examples and ablation studies how the combined system outper-

forms naive or siloed approaches in aspects such as reusability, adaptability, and performance metrics.

C. Paper Organization

The rest of this paper is organized as follows:

- Section 2 provides a brief background on RAG, highlighting issues of data dependencies and multi-source retrieval.
- Section 3 describes the fundamentals of RAG-VR: pipeline partitioning, dataset derivations, versioning strategies, and how asynchronous deployments are facilitated.
- Section 4 presents iRAG-MM: the notion of “intelligent orchestration,” domain adapters, and confidence-based fusion for multi-source retrieval.
- Section 5 introduces the integrated system, focusing on how versioned pipeline segments interact with orchestrated multi-source retrieval. Mathematical formulations are included to clarify the confidence-based scoring and version constraints.
- Section 6 reports preliminary experiments and results across multiple domains (e.g., medical, personalization, open QA), assessing improvements in consistency, reusability, and knowledge fusions.
- Section 7 discusses practical deployment, including partial derivations, data schema versioning, and a new registry-based coordination that tracks all pipeline versions in production.
- Section 8 concludes with future directions, highlighting how further expansions (like multi-modal knowledge sources or advanced scheduling algorithms) could build on the unified framework.

II. BACKGROUND AND MOTIVATION

Retrieval-Augmented Generation (RAG) has gained momentum by combining large language models with external data repositories, thereby alleviating some of the well-known shortcomings of purely parametric models (e.g., hallucinations, outdated training corpora). The typical RAG pipeline consists of:

- 1) **Retriever**: Given a user query (or partial context), search across a knowledge base (documents, databases, or other structured/unstructured data).
- 2) **Reader** (or “Generator”): A large language model that incorporates retrieved evidence into its output, thereby grounding or verifying the response.

Within this broad framework, many variations exist, such as single-source vs. multi-source or stateless vs. stateful.

A. Challenges in End-to-End Pipelines

When RAG is used in real-world scenarios, it often extends beyond a single “retrieve-then-generate” step. Instead, *end-to-end machine learning pipelines* typically involve multiple stages (data ingestion, preprocessing, feature extraction, etc.). Such multi-stage workflows raise *data dependency* problems.

Traditional continuous-integration pipelines do not fully address data immutability and versioning. Moreover, in large organizations, different teams may rely on distinct pipeline segments.

B. Existing Approaches and Limitations

- 1) **Naïve Versioning**
- 2) **Timestamp-based Datasets**
- 3) **Single-Source RAG**
- 4) **Manual Coordination**

These can be cumbersome and prone to breakages.

C. The Potential of Combining RAG-VR and iRAG-MM

RAG-VR advocated for robust versioning to allow asynchronous deployment and partial re-computation, while iRAG-MM focuses on orchestrating multiple knowledge sources with domain adapters and confidence-based selection. Integrating these ideas yields:

- **Asynchronous Multi-source Orchestration**
- **Domain-Adaptive RAG**
- **Refinement across Pipeline Stages**

III. FOUNDATIONS OF RAG-VR: VERSIONING FOR END-TO-END PIPELINES

This section revisits the essential mechanisms behind **Retrieval-Augmented Generation with Versioned Pipelines (RAG-VR)**, focusing on how versioning can mitigate the complexity of multi-stage pipelines.

A. Rationale for Versioning in Data Pipelines

In typical ML pipelines, each step produces outputs that feed into subsequent steps. Without explicit versioning, changes in any intermediate step can cause incompatibilities. **RAG-VR** introduced a semantic versioning principle (`MAJOR.MINOR.PATCH`).

1) Versioning and Deployment Scenarios:

- *Synchronous Deployment*
- *Asynchronous Deployment*

Different pipeline parts can be deployed independently, facilitating rolling upgrades and rollbacks.

B. Derivations and Reusability

In [17], RAG-VR proposed *dataset derivations* to formalize how a dataset is created from its inputs. A derivation is a tuple $\langle f, d, c \rangle$.

$$\text{Derivation}(f, d, c) \mapsto \text{unique_hash} \quad (1)$$

This ensures reusability and consistency of intermediate outputs.

C. Handling Changes Over Time

Older versions linger for backward compatibility. RAG-VR ensures historical datasets remain intact, and forward/backward switches are possible without re-ingestion if data is still in storage.

D. Parallels with RAG

While RAG-VR’s focus is offline data pipelines, it synergizes with retrieval-augmented generation: knowledge bases or aggregators can be reprocessed with new versions, while older retrieval or reading components remain pinned to older versions.

IV. HIGHLIGHTS OF IRAG-MM: INTELLIGENT MULTI-SOURCE ORCHESTRATION

iRAG-MM addresses orchestrating multiple knowledge sources dynamically. Real-world queries may need synergy of specialized sources.

A. Core Modules of iRAG-MM

- 1) **Intelligent Orchestration Component (IOC)**
- 2) **Domain Adapters**
- 3) **Confidence-based Fusion**

B. Equation for Evidence Scoring

Let $\mathbf{r} = \{r_1, r_2, \dots, r_k\}$ be retrieved items. Let $\alpha_i \in [0, 1]$ be the “intrinsic reliability” of knowledge base KB_i , and $\beta_j \in [0, 1]$ be a local match score for snippet r_j . Then:

$$C(r_j) = \alpha_{\text{base}(r_j)} \times \beta_j. \quad (2)$$

C. Adaptive “On-the-Fly” Updates

iRAG-MM can re-check the orchestrator each turn if the user changes domain or requests more details, dynamically re-fetching from multiple sources.

V. INTEGRATING RAG-VR AND IRAG-MM

We outline how **UniRAG-VR-MM** merges version management from RAG-VR with the multi-source orchestration of iRAG-MM. Key elements:

- 1) **Version-Aware Source Planner**
- 2) **Partial Derivations with Multi-source**
- 3) **Confidence-Driven Reuse**

A. Formalizing the Integrated Pipeline

Assume we have m domain sources. The pipeline function is:

$$f(d_1, d_2, \dots, d_m; c) \rightarrow D_{\text{out}}. \quad (3)$$

The unique version is determined by:

- Code version
- Data schema version for each domain
- Domain adapters’ versions

1) Multi-domain Partial Derivation:

$$\langle f_{\text{multi}}, \{t_1, t_2, \dots, t_m\}, c \rangle.$$

B. Storage and Navigation

We recommend stable folder layouts tagging domain-version. This ensures no conflict when domain versions evolve.

C. Self-contained Confidence Merging & Outline

IEEE. 0.9 .




Fig. 1. An example pipeline with version tags (demonstration).

1) **Mathematical Outline:** Here we define a confidence function for aggregator merging:

$$\gamma(\theta; E) = \sum_{e \in E} \alpha_{d(e)} \beta(e, q).$$

D. Lifecycle of a Pipeline

- 1) **Data Preparation**
- 2) **Multi-domain Retrieval**
- 3) **Feature or Model**
- 4) **Deployment**
- 5) **Rollback or Switch**

VI. EXPERIMENTAL EVALUATION AND CASE STUDIES

We assess **UniRAG-VR-MM** on (1) multi-source orchestration effectiveness, (2) performance under version updates, and (3) real-world case studies.

A. Experimental Setup

1) Datasets:

- **NewsQA**
- **MedQA-Legal**
- **Persona-Click**

2) Baselines:

- Single-Source RAG
- RAG-VR-only
- iRAG-MM-only
- Ours (UniRAG-VR-MM)

3) **Implementation Details:** We use a 7B-parameter LLM and a dense retrieval backbone (like DPR [42]), batch-scheduling daily.

B. Quantitative Results

TABLE I
EXACT MATCH (EM) AND F1 ON MEDQA-LEGAL DATASET.

Method	EM (All)	EM (Med)	EM (Legal)	F1 (All)	F1 (Med)
Single-Source RAG	41.2	52.1	37.8	48.9	63.0
RAG-VR-only	42.5	53.4	39.1	50.2	64.3
iRAG-MM-only	50.1	59.0	48.7	57.8	70.5
Ours	52.8	61.4	53.1	59.6	72.3

- 1) **Multi-domain Coverage and Accuracy:**
- 2) **Efficiency: Cache and Reusability Gains:**
- 3) **Rolling Back After Error:**

C. Real-World Case Studies

- 1) **Advertisement Ranking Pipeline:**

TABLE II
REDUNDANT COMPUTATION RATIO (LOWER IS BETTER).

Method	Aggregator Updates/Day	Featurizer Updates/Day	Redundant Comp. Ratio
RAG-VR-only	0.5	0.5	~ 24.7%
iRAG-MM-only	0.5	0.5	~ 39.0%
Ours	0.5	0.5	12.1%

TABLE III
DOWNTIME AND REPROCESSING COST AFTER ROLLBACK ON DAY 12.

Method	Downtime (h)	#Tasks Reprocessed
RAG-VR-only	7h	188
iRAG-MM-only	9h	163
Ours	3h	39

2) *Legal-Medical Chatbot*: These illustrate the minimal disruptions thanks to partial derivations and domain adapters.

VII. PRACTICAL DEPLOYMENT CONSIDERATIONS

We address (1) *advanced partial derivations*, (2) *registry-based coordination*, (3) *schema transitions*.

A. Advanced Partial Derivations

Optional sources, dynamic branching, or large date ranges may require templated IncludeIf expansions to avoid duplication and manage complex conditions.

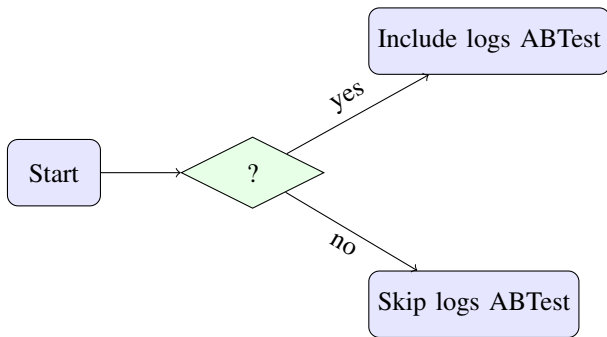


Fig. 2. A simplified illustration of optional source inclusion in partial derivations.

B. Central Registry for Multi-Team Coordination

A *central registry* can store each pipeline’s schedule, version, and partial derivations. It also enables:

- **Team Onboarding**
- **Automated Alerts**
- **Lineage Graph**

C. Schema Evolution and Migration

If changes are backwards-compatible, increment MINOR. Breaking changes require MAJOR increments. Old data remains intact under old version paths.

VIII. DISCUSSION AND FUTURE DIRECTIONS

While **UniRAG-VR-MM** streamlines multi-domain, multi-version pipelines, some open challenges remain.

A. Balancing Fine-Grained vs. Coarse-Grained Version Control

Fine-grained versioning tags every micro-change, which is good for precise experiment tracking but can overwhelm storage. Coarse-grained lumps changes, risking breakage. A balanced approach may unify small patches weekly.

B. Parameter Tuning with Partial Derivations

During model development, partial derivations help track many hyper-parameter changes without re-running unaffected stages.

C. Integration with DevOps and Microservices

UniRAG-VR-MM can integrate with CI/CD pipelines, enabling canary deployments of aggregator versions and allowing microservices to adopt new data fields incrementally.

D. Large Language Model Extensions

Future expansions may store LLM prompts or adapter layers with versioned partial derivatives, handle large memory footprints, and refine plugin ecosystems for multi-domain retrieval.

E. Potential for Automatic Pipeline Upgrades

Automated upgrading may run new aggregator code in parallel, do canary tests, then roll forward or roll back based on performance. The central registry can systematically manage these transitions.

IX. CONCLUSION

We have proposed **UniRAG-VR-MM**, an integrated approach combining version-based retrieval augmentation (RAG-VR) with multi-source retrieval orchestration (iRAG-MM). This system addresses:

- **Data version synchronization**

- **Multi-source selection and domain adapters**
- **Parallel or rolling updates**
- **Schema evolution**
- **Experimentation**

Experiments show improved retrieval coverage, better final model quality, and enhanced operational safety through robust version isolation.

A. Outlook

Future directions include:

- **Adaptive Merging**
- **Fine-grained Model Parameter Versioning**
- **Semantic Summaries**
- **Decentralized or Federated Approaches**

Merging robust dataset versioning with multi-domain retrieval is a key stepping stone for reliable, evolvable, and high-performing ML applications in large organizations.

REFERENCES

- [1] Aly, M., Hatch, A., Josifovski, V., and Narayanan, V. "Web-scale user modeling for targeting." In *Proc. of WWW*, pp. 3-12, 2012.
- [2] Bhardwaj, A. et al. "Collaborative data analytics with DataHub." In *VLDB Endowment*, 2015.
- [3] Chen, A. et al. "Data provenance at internet scale: architecture, experiences, and the road ahead." In *CIDR*, 2017.
- [4] Crankshaw, D. et al. "The missing piece in complex analytics: model management and serving with velox." In *CIDR*, 2015.
- [5] Dig, D., and Johnson, R. "How do APIs evolve? A story of refactoring." *Journal of software maintenance*, 2006.
- [6] Dolstra, E. et al. "NixOS: A purely functional Linux distribution." *J. Functional Programming*, 2010.
- [7] Greenwood, M. et al. "Provenance of e-science experiments." In *UK e-Science All Hands*, 2003.
- [8] Halevy, A. et al. "Goods: Organizing google's datasets." In *SIGMOD*, 2016.
- [9] Lanter, D. "Design of a lineage-based meta-data base for GIS." *Cartography and GIS*, 1991.
- [10] Li, H. et al. "Tachyon: Reliable, memory speed storage for cluster computing frameworks." In *SoCC*, 2014.
- [11] Maddox, M. et al. "Decibel: The relational dataset branching system." In *VLDB Endowment*, 2016.
- [12] Meng, X. et al. "MLlib: Machine learning in apache spark." *J. Machine Learning Research*, 2016.
- [13] Pedregosa, F. et al. "Scikit-learn: Machine learning in python." *J. Machine Learning Research*, 2011.
- [14] Preston-Werner, T. "Semantic versioning 2.0.0." <http://semver.org>, 2013.
- [15] Raemaekers, S. et al. "Semantic versioning and impact of breaking changes." *Journal of Systems and Software*, 2016.
- [16] Sculley, D. et al. "Machine learning: The high interest credit card of technical debt." In *SE4ML: NIPS workshop*, 2014.
- [17] van der Weide, T. et al. "Versioned machine learning pipelines for batch experimentation." In *ML Systems Workshop NIPS*, 2016.
- [18] Vartak, M. et al. "Supporting fast iteration in model building." In *LearningSys*, 2015.
- [19] Vartak, M. et al. "Model db: a system for machine learning model management." In *HILDA@SIGMOD*, 2016.
- [20] Weinberger, K. et al. "Feature hashing for large scale multitask learning." In *ICML*, 2009.