A Database-Independent LLM Framework for Real-Time Authorization in Retrieval-Augmented Generation

Halil Yesil[1], Sumeyye Gultekin[2], Fatma Bozyigit[3], Baris Tekin Tezel[1], Moharram Challenger[3]
[1] **Dokuz Eylul University, Izmir, Turkey**
[2] **Cankaya University, Ankara, Turkey**
[3] **University of Antwerp and Flanders Make, Antwerp, Belgium**
halilyesilceng@gmail.com, smygltkn32@gmail.com
fatma.bozyigit@uantwerpen.be, baris.tezel@deu.edu.tr,
moharram.challenger@uantwerpen.be

Abstract

Retrieval-Augmented Generation (RAG) systems have become effective tools for improving language model capabilities by integrating external knowledge. However, these systems face significant security problems when handling sensitive information, particularly concerning unauthorized data access. This paper presents a new access-controlled RAG framework that includes dynamic security measures without needing manual database construction or extensive preprocessing of user-specific document collections. Our approach uses three specific knowledge bases and advanced language models to make real-time access control decisions. Experimental testing on a dataset of 167 posters from six organizations—University of Antwerp, Flanders Make, University of Ghent, University of Hasselt, KU Leuven, and VUB—shows 93% overall accuracy in access control decisions, with 92% accuracy for text-based control and 78% for image-based control. The framework enables quick deployment of secure RAG systems while maintaining efficiency and responsiveness.

Introduction

Retrieval-Augmented Generation (RAG) systems have become an effective method for improving language model abilities by incorporating external knowledge sources into the generation process (Lewis et al. 2020). These systems combine what large language models have learned (Brown et al. 2020) with information pulled from external documents. This approach helps them provide more accurate, up-to-date, and contextually relevant responses. However, using external data sources raises significant security and privacy issues that remain largely unaddressed in current systems (Carlini et al. 2021; Mireshghallah et al. 2022).

A major concern is the risk of accidental data leakage when RAG systems access and use sensitive information during retrieval and generation. Without appropriate safeguards, these systems might unintentionally expose confidential documents, personal details, or proprietary information to unauthorized individuals. This risk is particularly concerning in enterprise environments, healthcare, legal practices, and other fields where strict data management and access control rules are essential.

Standard access control methods are well-known in database systems and file management (Sandhu et al. 1996; Ferraiolo et al. 2001). However, they encounter unique challenges when applied to RAG systems. The fast nature of retrieval-augmented generation, where relevant documents are identified and processed in real-time based on user queries, requires access control decisions to be made at multiple points in the pipeline. The retrieval step decides which documents can be accessed, while the generation step determines how the retrieved information is included in responses. Both steps must align with user authorization levels and organizational access rules.

Current methods for securing RAG systems often rely on processing entire document collections in advance to create user-specific databases or on applying broad access controls at the system level (Chen 2024; Jo and Yang 2024). However, these approaches can be expensive in terms of computation and time, especially in settings with large document collections and diverse user groups. The effort needed to create individual knowledge bases for each user and document leads to significant scalability problems that limit the practical use of secure RAG systems.

This research addresses these issues by proposing a combined access control framework for RAG systems that can enforce authorization policies in real-time without extensive pre-processing of document collections. The proposed solution seeks to strike a balance between security requirements and system efficiency, ensuring that access control does not slow down the speed and scalability that make RAG systems attractive for practical application.

Related Work

The integration of access control mechanisms with Retrieval-Augmented Generation systems has become increasingly important as organizations strive to implement RAG systems while ensuring data security and privacy. Researchers have suggested various methods to tackle the challenges of creating secure RAG systems across different fields.

Chen (Chen 2024) puts forward a foundational approach that involves placing an access control layer between the retriever and generator components of a standard RAG pipeline. The study sets up an access control database that maintains clear mappings between users and their authorized documents. Before retrieving any document, the system checks permissions to confirm that users cannot access unauthorized content. Although this method is simple and effective, it requires manually defining all access rules, which may hinder scalability in rapidly changing environments.

Jo and Yang (Jo and Yang 2024) created an on-premise RAG system tailored for enterprise settings with strict data privacy needs. Their framework uses a relational database to hold detailed document access policies, which include department and role-based restrictions. The access control mechanism functions between the retrieval and generation phases, sorting documents based on user attributes before document retrieval starts. This method shows better scalability through dynamic access policies that adjust to shifting organizational structures and needs.

In healthcare, Chen et al. (Chen, Li, and Wang 2025) developed a role-based access control (RBAC) model to manage sensitive patient profiles in medical environments. Their system ensures that healthcare professionals, such as doctors and nurses, can retrieve data only in line with their specific authorization levels. This implementation stands out for meeting strict healthcare compliance standards, showcasing that RAG systems can be safely used in high-stakes clinical settings where data privacy is critical.

Namer et al. (Namer, Schmidt, and Rodriguez 2024) propose a complex two-layered filtering method that makes use of both metadata and graph-based structures for access control. Their approach first gathers extensive metadata from documents, including author details, departmental classification, and security levels. This metadata is then utilized to build a graph database where nodes represent users and documents, and edges illustrate access relationships. Before fetching documents, the system checks access rights through graph traversal, allowing intricate reasoning over permissions and supporting more detailed access control scenarios.

The most sophisticated approach found in current research comes from Jayasundara et al. (Jayasundara et al. 2024), who introduce RAGent. This framework automatically generates access control policies from natural language documents. Their six-stage pipeline includes coreference resolution to clarify ambiguous pronouns, Natural Language Access Control Policy (NLACP) extraction to find permission statements, context matching with existing data, embedding creation using LLaMA for semantic content mapping, synonym resolution for consistent terminology, and manual verification with iterative refinement. This method significantly cuts down the manual work involved in creating access control policies.

Despite the progress made in merging access control with RAG systems, some limitations remain. Most current methods demand considerable pre-processing of document collections or manual policy setting, which can create scalability issues in large deployments. Moreover, the computational demands of access control mechanisms often affect system performance, potentially restricting the practical use of these solutions.

Table 1 provides a comparative overview of the existing approaches discussed in the literature.

Table 1: Comparison of Access Control Approaches in RAG Systems

Approach	Year	DB Dep.	Data Type
Chen (Chen 2024)	2024	Yes	Textual
Jo & Yang (Jo and Yang 2024)	2024	Yes	Textual
Chen et al. (Chen, Li, and Wang 2025)	2025	Yes	Textual
Namer et al. (Namer, Schmidt, and Rodriguez 2024)	2024	Yes	Textual
Jayasundara et al. (Jayasundara et al. 2024)	2024	No	Textual

State of The Technology

The field of access control technology has changed significantly to meet the varying security needs of today's information systems. Knowing the current status of access control methods is crucial for creating strong security frameworks for RAG systems. Each method has its own advantages and limitations based on the context of deployment and security needs.

Access Control Models

Role-Based Access Control (RBAC) is one of the most widely used access control methods in businesses. In RBAC systems, users are assigned to specific roles such as "Administrator," "Manager," or "Employee." Permissions are given to these roles as a group, not individual users. This method simplifies permission management by cutting down on the work involved in managing permissions for each user. However, since RBAC depends on fixed role definitions, it lacks flexibility in situations where users need specific or context-based permissions. This makes it less effective in changing environments where access needs frequently shift.

Attribute-Based Access Control (ABAC) and **Policy-Based Access Control (PBAC)** provide more advanced access control by making authorization decisions based on a wide range of attributes linked to users, resources, and the environment. These attributes can include user roles, document classifications, geographic locations, time limits, and security clearance levels. While ABAC and PBAC systems offer detailed, context-aware access control that meets complex organizational needs, they demand thorough policy definition and complex evaluation engines. This leads to greater implementation complexity and increased computational demands.

Relationship-Based Access Control (ReBAC) sets access permissions based on the relationships between entities in a system. For instance, a user may get access to a

document if they are labeled as the "manager of" the document owner or are a "member of" a specific project team. This model works well in social networks and graph-based systems where relationships clearly define access limits. ReBAC also allows for dynamic access control by automatically changing permissions as relationships shift within the organization. This makes it ideal for collaborative environments with changing team structures.

Access Control Lists (ACL) provide document-level access control by keeping clear lists of users or groups with specific permissions for each resource. This method allows for precise control over individual documents and enables detailed permission management. However, ACLs become harder to scale and maintain in large document repositories because each resource needs individual permission management. The administrative task of managing ACLs increases significantly with the number of documents and users in the system.

Context-Based Access Control (CBAC) is an emerging model that applies to RAG systems. In this model, access decisions are based on the context of user queries and the sensitivity of the requested information. CBAC systems look at factors like content sensitivity, user intent, query patterns, and environmental conditions to make access control decisions. This method is particularly useful in RAG systems, where mixed-sensitivity content can be retrieved. It allows for real-time filtering to stop unauthorized information from being disclosed during retrieval and generation phases.

Fine-Grained Authorization (FGA) has become popular because of systems inspired by Google's Zanzibar. These systems use authorization rules like "user A can perform action B on resource C if condition D holds." Modern FGA frameworks, like OpenFGA and Okta FGA, offer scalable authorization systems. They can check permissions in real-time across complex organizational structures. These systems work well for RAG applications because they can evaluate permissions at query time without needing extensive preprocessing of user-specific datasets.

Comparative Analysis

Table 2 provides a comprehensive comparison of the various access control models, highlighting their key characteristics and applicability to different system architectures.

Technology Integration Challenges

The integration of traditional access control models with RAG systems poses unique challenges that need careful thought during system design. The changing nature of information retrieval and generation in RAG systems requires access control methods that can work effectively at query time without slowing down system performance. Traditional models like RBAC and ACL, even though they are well-established, might not offer the flexibility and scalability needed for RAG deployments in various organizational settings.

Modern methods like CBAC and FGA show promise for RAG applications because they can make context-based authorization decisions in real-time. However, to implement these advanced access control models successfully, it is necessary to have sophisticated policy engines and to carefully consider the computational resources required for complex authorization evaluations.

Table 2: Comparison of Access Control Models

Model	Summary
Role-Based Access Control (RBAC)	Assigns permissions to roles, not individuals. It is simple but inflexible for changing needs
Attribute-/Policy-Based Access Control (ABAC/PBAC)	Grants access based on user and resource attributes. It is flexible but complex to manage and implement
Relationship-Based Access Control (ReBAC)	Uses relationships, such as "manager of," for dynamic access control in collaborative settings
Access Control Lists (ACL)	Lists specific users or groups for each resource. This approach is precise but hard to scale across large datasets
Context-Based Access Control (CBAC)	Takes into account the query context, like sensitivity or user intent. It works well for dynamic RAG systems
Fine-Grained Authorization (FGA)	Uses tuple-based rules that allow scalable, real-time permission checks. This is suitable for modern RAG architectures

Dataset Preparation

Project posters that are publicly available were collected from various websites. While the raw dataset of posters is accessible to everyone, we processed the data to create an access control mechanism for our experimental setup. Each poster includes information about the organization that owns it. The dataset has 167 posters from six different organizations: University of Antwerp, Flanders Make, University of Ghent, University of Hasselt, KU Leuven, and VUB. This offers a diverse foundation for evaluating access control mechanisms in a multi-organizational context. The distribution of posters among the six organizations is shown in Figure 1.

Poster Preparation

Raw posters are available to the public, and their ownership can be identified from the poster text. To create a realistic access control scenario, we changed the original public text of the posters to include various security classifications like confidential, copyright, and private. However, some posters were intentionally kept with their original public status to ensure diversity in access levels. This change allows us to apply access control mechanisms across different security classifications.

We also added organization logos and security level indicators (confidential, private, and copyright) to the posters as visual elements. These additions have two main benefits. They provide a clear visual identification of the owning organization and indicate the security level of each poster. This

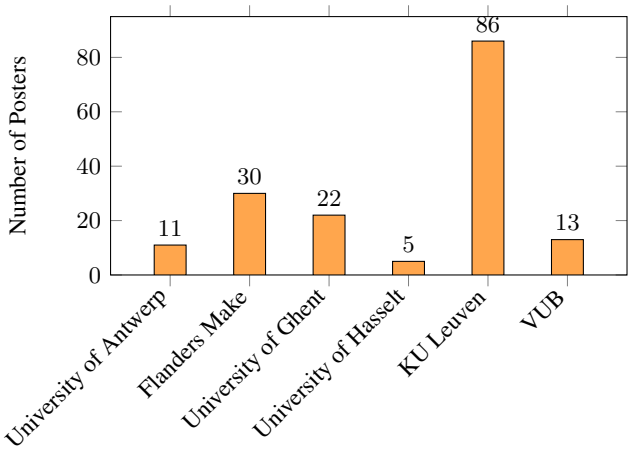

Figure 1: Number of Posters per Organization

helps both people and automated systems verify access control.

Technical Implementation

The dataset preparation pipeline used Python libraries for efficient document processing and manipulation. PyMuPDF (fitz) was the main library for working with PDFs. It allowed access to content, extraction of metadata, and the addition of text-based security classifications. OpenCV handled image processing tasks, such as resizing logos, positioning, and overlaying images while keeping the visual quality and integrity of the documents intact.

The mxbai-embed-large model produced 1024-dimensional vector embeddings of document chunks. These embeddings captured semantic relationships, allowing for efficient similarity searches during retrieval. The qwen2.5:72b and qwen2.5vl:72b models were accessed through the Ollama framework API for local deployment and inference. They supported both text-based analysis and vision-language tasks to extract organizational and security information from visuals. This combined technical setup creates a strong and scalable preprocessing pipeline that can manage various document types while staying efficient.

Hybrid Security Level Annotation To evaluate the strength of the access control systems for different document types, we used a mixed annotation strategy for classifying security levels. The dataset was organized into three separate categories of posters based on security level indicators:

Text-only security indicators: Some posters include security level classifications, such as confidential, private, copyright, and public, that are only found within the text of the document. These posters do not have visual security badges or logos. This means the access control system must extract and understand security classifications using only the natural language in the document text.

Image-only security indicators: Another subset uses visual elements to convey security level information. It features logo badges that represent confidential, private, and copyright classifications. These posters do not include clear textual security level statements. This puts the system's ability to analyze visual content and the image access control knowledge base to the test.

Hybrid text-image security indicators: The third category includes posters that show security level information in both text and images. These documents feature clear security level text along with matching visual logo badges. This allows for an assessment of the system's capacity to combine information from different formats and check for consistency among various sources.

This hybrid approach to security level annotation allows for a thorough evaluation of the system's multimodal access control features. It tests the strength of both text-based and image-based security classification methods, both separately and together.

Hybrid Organization Level Annotation Similarly to the security level annotation strategy, we added organizational affiliation information to the data set using a mixed approach. This helped us assess the system's ability to recognize organizational ownership in various content types. The organizational annotation strategy includes two main categories:

Text-only organization indicators: Many posters have details about organizational affiliation only in the text. These documents clearly mention the owning organization using institutional names without showing any visual logos.

Hybrid text-logo organization indicators: The other posters include organizational information with written references and visual logos. These documents show institutional branding elements and written identifiers, allowing the system to verify organizational affiliation through different information channels.

The careful design of this hybrid annotation framework meets several evaluation goals. First, it allows separate assessments of text-based and image-based ways to identify organizations. Second, it helps evaluate how well the system deals with real-world situations, where organizational information might come in different forms depending on the type of document and source. Third, it sheds light on how text and visual access controls work together. This shows whether using multiple modes improves overall system accuracy and reliability.

Knowledge Base Generation

In this study, we created three different knowledge bases to support various aspects of the access control system. The first is the poster content knowledge base, which was developed to support the RAG mechanism. This knowledge base allows for answer generation based on relevant pieces that match user queries. It provides the essential retrieval capability for the system. The structure of the poster content knowledge base is shown in Figure 2.

The second knowledge base is the text access knowledge base. It has two access levels: organization-level and security-level permissions. Organization-level access control sets permissions based on organizational affiliation. Security-level access control manages permissions according to document classification.

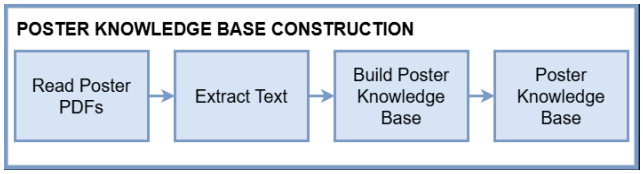

Figure 2: Poster Knowledge Base

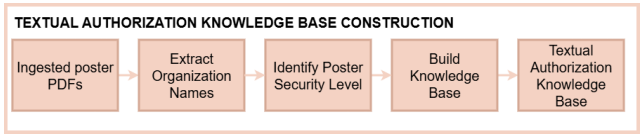

Figure 3: Text Access Knowledge Base

The third knowledge base is the image access knowledge base. It was created using the same method as the text access knowledge base. For generating the image access control knowledge base, qwen2.5vl:72b is used to process visual content and gather information on organizational and security levels. This knowledge base offers details about organizational ownership and security level classifications for visual content. It ensures thorough access control across various media types in the poster dataset.

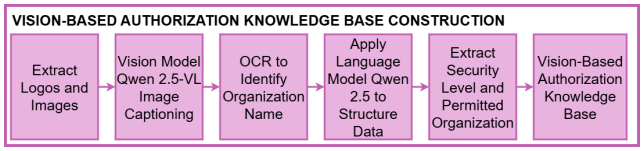

Figure 4: Image Access Control Knowledge Base

Proposed Methodology

This research presents a new access-controlled Retrieval-Augmented Generation (RAG) framework that adds security measures to the traditional RAG process. The approach includes three main modules that operate in sequence to guarantee secure information retrieval and generation: chunk retrieval, access control verification, and controlled answer generation. The system ensures that users can only access information from documents they are allowed to view. It also maintains the effectiveness of the RAG mechanism for retrieving knowledge and generating answers. You can see the overall design of the proposed system in Figure 5.

The proposed framework tackles the crucial problem of stopping unauthorized access to sensitive information in RAG systems by using detailed access control for both text and visual content. This method uses multiple knowledge bases and advanced language models to make real-time access control decisions without needing extensive preprocessing for individual users.

Retrieve Similar Chunks

The initial phase of the methodology involves retrieving relevant document chunks based on user queries. When a user submits a query, the system first converts it into a high-dimensional vector representation using the mxbai-embed-large embedding model. This model was chosen for its strong ability to capture semantic relationships and contextual meaning in text.

After creating the query embedding, the system performs a similarity search against the poster content knowledge base using cosine similarity metrics. The knowledge base contains pre-computed embeddings of document chunks from the poster dataset, which allows for efficient matching between user queries and relevant content. The similarity search algorithm finds and retrieves the five chunks that are most similar to the user query, along with their corresponding PDF document names and metadata.

The retrieval process uses a vector database structure, which allows for quick similarity calculations across large document collections. Each retrieved chunk includes contextual information, such as the document source, chunk position within the document, and relevance scores. This information is important for access control verification and ensures traceability throughout the process. The chunk retrieval process is shown in Figure 6.

Access Control

The access control module is the main security part of the proposed method. It uses a dual-layer verification system that checks both text and visual access permissions. When the system gets the PDF names from the similarity search, it queries both the text and image access control databases to gather detailed access control information for each document.

The access control decision process uses the user's email address as the main identifier to determine their organization. The system extracts the organization domain from the user's email and matches it to the right organization in the access control framework. This method allows for user identification without needing manual user registration or role assignments.

For each retrieved document, the system uses the qwen2.5:72b large language model to analyze the overlap between user permissions and document requirements. The model processes three key inputs: the user's organization based on their email address, the document's ownership information from the access control databases, and the document's security classification level. The language model uses contextual reasoning to assess whether the user meets both organizational and security requirements to access each specific document.

The access control evaluation examines different permission aspects, such as organizational membership and security clearance levels. The qwen2.5:72b model's reasoning skills help it handle complicated access control situations, including documents involving various organizational stakeholders or different security classifications.

RAG Mechanism

The final phase carries out the controlled generation process. It creates answers using only the information from documents that passed access control verification. The system

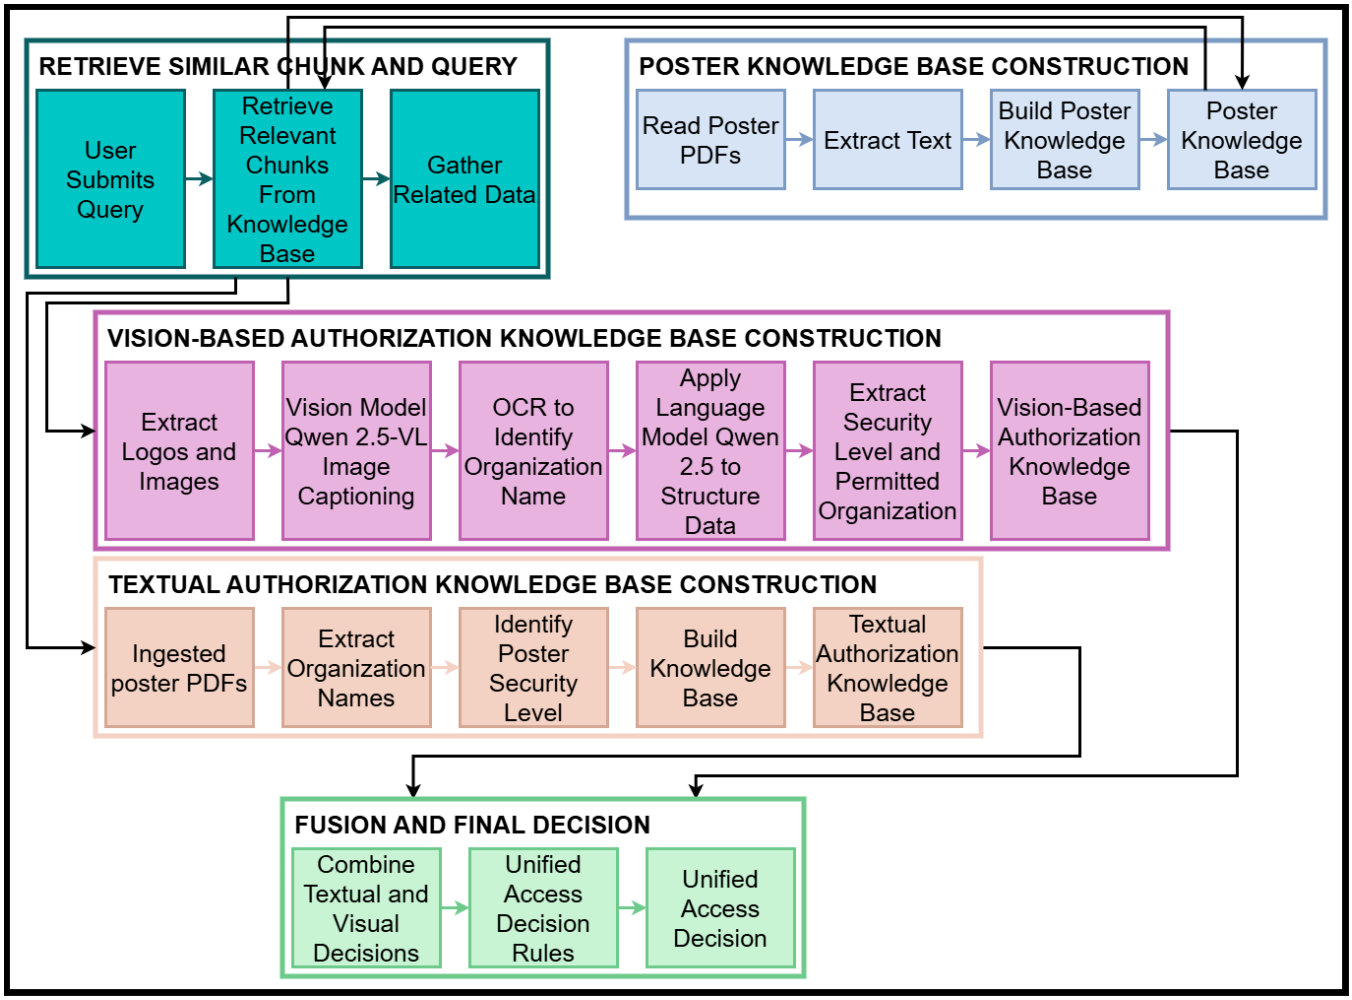

Figure 5: Access Control with Retrieval-Augmented Generation Architecture

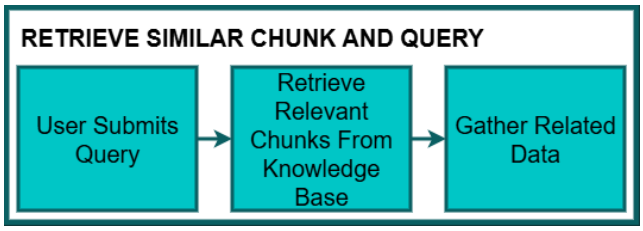

Figure 6: Retrieve Similar Chunk

takes the filtered set of available chunks identified by the access control module and combines them with the user's original query to generate a complete response.

During this process, the qwen2.5:72b language model is used differently than in the access control phase. Here, it focuses on synthesizing information and generating natural language instead of making access decisions. The model receives the user's original query along with the combined text from all accessible document chunks. This ensures that the generated response relies solely on approved information sources. The fusion process and final decision generation are shown in Figure 7.

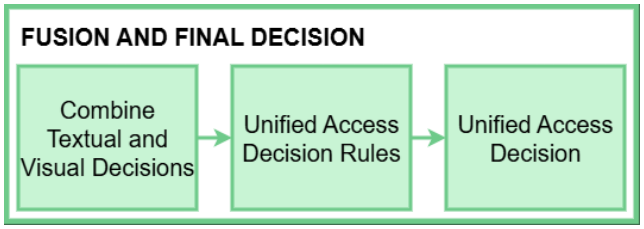

Figure 7: Fusion and Final Decision

Evaluation

The evaluation of the proposed access-controlled RAG system focuses on assessing the accuracy and effectiveness of the access control mechanisms implemented within the framework. A comprehensive evaluation methodology was designed to measure the system's ability to correctly iden-

tify and enforce access permissions across different content types and organizational boundaries.

Evaluation Methodology

The evaluation process focuses on assessing the accuracy of the access control knowledge bases. These bases are essential for security enforcement in the proposed system. The evaluation framework uses a binary classification method, where access control decisions are classified as either correct or incorrect based on ground truth annotations.

For every document in the evaluation dataset, ground truth labels were created by manually marking the allowed organizations and security levels based on the document's content and built-in security indicators. These markings serve as the standard to compare the system's automated access control decisions against.

Multimodal Evaluation Strategy The evaluation method was created to measure the system's performance within the hybrid annotation framework used during dataset preparation. The evaluation process includes several assessment dimensions that relate to the different modality configurations found in the dataset.

Security level access control evaluation: For security level classification, the evaluation framework measures accuracy across three different poster categories: posters with text-only security indicators, posters with image-only security indicators (logo badges), and posters with a mix of text and image security indicators. This structured evaluation helps quantify the performance of the text access control knowledge base, the image access control knowledge base, and the combined multimodal system. The evaluation metrics track the system's ability to accurately identify confidential, private, copyright, and public classifications in all modality configurations.

Organization level access control evaluation: Organizational affiliation identification is assessed in two main categories: posters with text-only organizational indicators and posters that combine text and logos. The evaluation looks at the system's ability to extract and match organizational affiliations from six organizations in the dataset: University of Antwerp, Flanders Make, University of Ghent, University of Hasselt, KU Leuven, and VUB. This is done regardless of whether the organizational information appears as text, logos, or a mix of both.

Integrated multimodal performance assessment: Beyond independent evaluation of text-based and image-based access control methods, the evaluation framework looks at the performance of the combined system when merging results from both types. This assessment checks if combining text and visual access control decisions improves overall accuracy compared to using just one method. The evaluation studies the combined effect of using both types of information, where both text and image analysis play a role in the final access control decision.

The evaluation strategy allows for a detailed understanding of how the system performs under different document setups. It offers insights into the strengths and weaknesses of each access control method and their combined effects in a multimodal framework.

Results and Analysis

The experimental evaluation shows how effective the proposed access control mechanisms are, with different performance levels seen across various content types. The text access control module reached an accuracy of 92%, which means it reliably extracts and classifies organizational ownership and security levels from the poster documents.

The image access control module achieved an accuracy of 78%. This reflects the challenges that come with analyzing visual content and the difficulty of extracting meaningful information from logos and security level indicators found in poster images. Even though this performance is lower than the text-based analysis, it is still an important step forward in automated visual access control for document security systems.

When the outputs of both the text and image access control modules were combined, the system as a whole achieved an accuracy of 93%. This improvement over the text-only performance suggests that using both methods together provides helpful information that strengthens access control decisions.

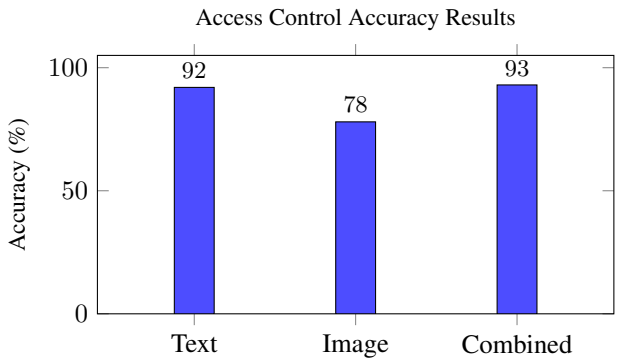

Figure 8: Access Control Accuracy Results (Bar Chart)

Conclusion

This research presents a new way to integrate access control mechanisms into Retrieval-Augmented Generation systems. It tackles the important issue of preventing unauthorized access to sensitive information in enterprise RAG deployments. The proposed framework shows that effective access control can be applied without losing the efficiency and responsiveness that make RAG systems appealing for practical use.

The methodology introduced in this study has clear advantages over current methods. It removes the need for extensive preprocessing of document collections for each user. By implementing three specialized knowledge bases and strategically using advanced language models, the system makes real-time access control decisions. This addresses a major scalability issue highlighted in existing literature.

The experimental evaluation confirms the effectiveness of the proposed approach, reaching an overall accuracy of 93% in access control decisions across a diverse dataset of 167 posters from six different organizations: University of Antwerp, Flanders Make, University of Ghent, University of Hasselt, KU Leuven, and VUB. The multimodal access control mechanism combines both textual and visual content analysis and performs better than single-modality approaches.

One of the key contributions of this work is showing that it is possible to quickly set up secure RAG systems without sacrificing security effectiveness. The framework can automatically extract organizational affiliations from user email addresses and dynamically assess access permissions. This eliminates the slow manual configuration processes found in current solutions.

Future research should explore integrating temporal access controls, where permissions may change over time. It should also look into creating more detailed access control policies that consider factors beyond organizational affiliation and security levels. This can pave the way for more trustworthy Agents and Multi-agent Systems (Demirkol et al. 2011; Kardas et al. 2012; Demirkol et al. 2012). Additionally, testing scalability with larger document collections and user groups would provide valuable insights into the system's performance in real-world enterprise settings.

References

Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 33: 1877–1901.

Carlini, N.; Tramer, F.; Wallace, E.; Jagielski, M.; Herbert-Voss, A.; Lee, K.; Roberts, A.; Brown, T.; Song, D.; Erlingsson, Ú.; Oprea, A.; and Raffel, C. 2021. Extracting Training Data from Large Language Models. In *30th USENIX Security Symposium*, 2633–2650.

Chen, L. 2024. An exploration of integrating access control in retrieval-augmented generation systems. *Journal of Information Security*, 15(3): 45–62.

Chen, M.; Li, X.; and Wang, Y. 2025. Integrating role-based access control in healthcare RAG systems. *IEEE Transactions on Biomedical Engineering*, 72(1): 234–245.

Demirkol, S.; Challenger, M.; Getir, S.; Kosar, T.; Kardas, G.; and Mernik, M. 2012. SEA_L: a domain-specific language for semantic web enabled multi-agent systems. In *2012 Federated Conference on Computer Science and Information Systems (FedCSIS)*, 1373–1380. Ieee.

Demirkol, S.; Getir, S.; Challenger, M.; and Kardas, G. 2011. Development of an agent based e-barter system. In *2011 International Symposium on Innovations in Intelligent Systems and Applications*, 193–198. IEEE.

Ferraiolo, D. F.; Sandhu, R.; Gavrila, S.; Kuhn, D. R.; and Chandramouli, R. 2001. Proposed NIST Standard for Role-Based Access Control. *ACM Transactions on Information and System Security*, 4(3): 224–274.

Jayasundara, P.; Thompson, D.; Kumar, R.; and Patel, S. 2024. RAGent: Automated access control policy generation for retrieval-augmented systems. In *Proceedings of the ACM Conference on Computer and Communications Security*, 789–802.

Jo, S.; and Yang, K. 2024. A study on secure on-premise RAG systems for enterprise environments. In *Proceedings of the International Conference on Enterprise Security*, 123–135.

Kardas, G.; Challenger, M.; Yildirim, S.; and Yamuc, A. 2012. Design and implementation of a multiagent stock trading system. *Software: Practice and Experience*, 42(10): 1247–1273.

Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; Riedel, S.; and Kiela, D. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems*, volume 33, 9459–9474.

Mireshghallah, F.; Goyal, K.; Uniyal, A.; Berg-Kirkpatrick, T.; and Shokri, R. 2022. Quantifying Privacy Risks of Masked Language Models Using Membership Inference Attacks. *arXiv preprint arXiv:2203.03929*.

Namer, A.; Schmidt, B.; and Rodriguez, C. 2024. Retrieval-augmented generation with graph-based access control. In *Advances in Neural Information Processing Systems*, volume 37, 1567–1580.

Sandhu, R. S.; Coyne, E. J.; Feinstein, H. L.; and Youman, C. E. 1996. Role-Based Access Control Models. *Computer*, 29(2): 38–47.