

Lost in OCR Translation? Vision-Based Approaches to Robust Document Retrieval

Alexander Most; Joseph Winjum; Manish Bhattarai;
Shawn Jones; Nishath Rajiv Ranasinghe; Ayan Biswas;
Dan O'Malley

Accessibility Disclaimer:

For a more accessible version of this document, please submit an accessibility request form through the Montana State University Library website.

Lost in OCR Translation? Vision-Based Approaches to Robust Document Retrieval

Alexander Most
AMost@lanl.gov
Los Alamos National Laboratory
Los Alamos, NM, USA

Joseph Winjum
Joseph.Winjum@gmail.com
Gianforte School of Computing
Montana State University
Bozeman, MT, USA

Ayan Biswas
Ayan@lanl.gov
Los Alamos National Laboratory
Los Alamos, NM, USA

Shawn Jones
SMJones@lanl.gov
Los Alamos National Laboratory
Los Alamos, NM, USA

Nishath Rajiv Ranasinghe
Ranasinghe@lanl.gov
Los Alamos National Laboratory
Los Alamos, NM, USA

Dan O'Malley
OMalleD@lanl.gov
Los Alamos National Laboratory
Los Alamos, NM, USA

Manish Bhattarai
Ceodsppectrum@lanl.gov
Los Alamos National Laboratory
Los Alamos, NM, USA

ABSTRACT

Retrieval-Augmented Generation (RAG) has become a popular technique for enhancing the reliability and utility of Large Language Models (LLMs) by grounding responses in external documents. Traditional RAG systems rely on Optical Character Recognition (OCR) to first process scanned documents into text. However, even state-of-the-art OCRs can introduce errors, especially in degraded or complex documents. Recent vision-language approaches, such as ColPali, propose direct visual embedding of documents, eliminating the need for OCR. This study presents a systematic comparison between a vision-based RAG system (ColPali) and more traditional OCR-based pipelines utilizing Llama 3.2 (90B) and Nougat OCR across varying document qualities. Beyond conventional retrieval accuracy metrics, we introduce a semantic answer evaluation benchmark to assess end-to-end question-answering performance. Our findings indicate that while vision-based RAG performs well on documents it has been fine-tuned on, OCR-based RAG is better able to generalize to unseen documents of varying quality. We highlight the key trade-offs between computational efficiency and semantic accuracy, offering practical guidance for RAG practitioners in selecting between OCR-dependent and vision-based document retrieval systems in production environments.

ACM Reference Format:

Alexander Most, Joseph Winjum, Ayan Biswas, Shawn Jones, Nishath Rajiv Ranasinghe, Dan O'Malley, and Manish Bhattarai. 2025. Lost in OCR Translation? Vision-Based Approaches to Robust Document Retrieval. In *ACM Symposium on Document Engineering 2025 (DocEng '25)*, September 2–5, 2025, Nottingham, United Kingdom. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3704268.3742698>

1 INTRODUCTION

Large Language Models (LLMs) have shown improvements across the landscape of natural language processing, enabling promising performance increases across a multitude of tasks such as semantic analysis, question answering, machine translation, and text descriptions. However, despite their strengths, these models suffer from fundamental limitations due to their reliance on static training corpora, often leading to hallucinations or outdated information. Retrieval-Augmented Generation (RAG) addresses this limitation by allowing LLMs to retrieve and cite information from external sources during inference [10]. In a typical RAG pipeline, user queries are transformed and stored as vector embeddings. The most relevant documents are then passed into the LLM's context window along with the original query to guide generation. Leveraging this approach, RAG has been leveraged to boost code translation quality [4][5] and has likewise been shown to curb hallucinations [3].

In addition to improving response accuracy, the retrieval component of RAG substantially facilitates the application of LLMs to proprietary and evolving document collections. Unlike traditional fine-tuning methods—which are costly, time-consuming, and must be repeated whenever new documents are introduced—RAG enables real-time integration of the latest information through on-demand retrieval. Consequently, RAG has been widely adopted in commercial production settings, powering proprietary systems like ChatGPT, Claude, Perplexity, and Gemini, although detailed technical implementations remain largely unpublished in peer-reviewed venues.

This work specifically addresses the application of RAG to scanned or digitized documents. Traditionally, handling such documents involves first extracting text using OCR algorithms prior to embedding and retrieval. However, OCR preprocessing frequently introduces significant noise, especially in scenarios involving low-quality images, handwritten annotations, or complex layouts. Recent VLMs, such as ColPali [6], offer a promising alternative by



directly embedding document images into unified multimodal vector spaces without OCR. These vision-based methods inherently preserve critical spatial relationships, formatting cues, and visual nuances that OCR processes often overlook or misinterpret.

Despite the potential advantages of direct image embedding, critical gaps persist in the current literature. Existing benchmarks for document retrieval systems, such as ViDoRe [6], predominantly evaluate performance using clean, high-quality documents, conditions rarely representative of practical real-world scenarios. Additionally, prior comparative evaluations—including those involving ColPali—have often been limited by comparisons against outdated OCR baselines like Tesseract [14], failing to represent the true capabilities of modern OCR-based systems and potentially inflating the perceived advantages of VLM-based methods. To date, the field lacks rigorous comparative assessments of OCR-to-RAG versus VLM-to-RAG pipelines across varied document quality conditions, particularly emphasizing downstream semantic accuracy and practical robustness.

Moreover, existing vision-language models suitable for direct embedding (e.g., ColPali, ColQwen, ColSmol) typically employ relatively lightweight architectures with parameter counts ranging between 2–7 billion, whereas current OCR-to-RAG pipelines can leverage significantly larger language models. Specifically, in this study, our OCR-based pipeline integrates Llama 3.2 (90 billion parameters), potentially affording substantial advantages in downstream semantic quality.

While one could theoretically adapt a 90B parameter model for the VLM role in ColPali, implementing this architecture presents significant practical challenges related to performance and memory usage. ColPali relies on generating and storing multiple dense vector embeddings per document (representing image patches), and using a VLM of 90B magnitude would dramatically increase the computational time needed to generate these numerous embeddings. Additionally, the memory and disk space required to store these multi-vector representations would likely become prohibitive for substantial document collections. The subsequent question embedding before the retrieval step would also face substantial slowdowns, potentially rendering the RAG system unusable for interactive applications. Although an operation might tolerate the costs associated with document embedding, the degraded query response times would likely make the system unacceptable to end users.

Our empirical results show that even when using the same query encoder (Qwen2 7B), the VLM-based pipeline exhibits slower query times than the OCR-based system. This slowdown stems from the need to compare the query embedding against many patch-level image embeddings per document during late-interaction retrieval, making latency a practical bottleneck even with mid-sized models.

Given the increase in retrieval speed inherent in scaling the VLM component within this specific dense retrieval architecture, this study focuses on comparing representative, readily deployable model sizes characteristic of each paradigm (lightweight VLM vs. OCR + large LLM). This allows for an evaluation of the trade-offs presented by systems configurations commonly encountered in practice.

This study rigorously compares a VLM-based RAG pipeline (ColPali) against a competitive OCR-based pipeline utilizing Llama 3.2

OCR. We systematically evaluate their performance across multiple document quality levels, explicitly incorporating visual degradation and complexity reflective of real-world conditions. Furthermore, we introduce a comprehensive semantic answer evaluation benchmark that extends beyond traditional retrieval metrics, enabling a deeper understanding of each pipeline's practical strengths and limitations.

We initially anticipated that the integrated visual-textual embedding of VLM-based approaches would offer superior robustness to visual noise, layout variations, and OCR-induced inaccuracies, potentially resulting in higher reliability for practical applications involving visually imperfect documents. However, our empirical findings indicate that VLM-based approaches struggle to generalize effectively to unseen data that they have not been explicitly fine-tuned on.

1.1 Contributions

This study advances the state-of-the-art (SOTA) through three core contributions: (1) We present the first systematic empirical comparison between two leading paradigms for document retrieval augmentation: VLM-based RAG system, specifically ColQwen2 (7 billion parameters), and an advanced OCR-based RAG pipeline leveraging Llama 3.2 (90 billion parameters). To the best of our knowledge, previous evaluations have not explicitly addressed the substantial model-size disparity between lightweight VLM architectures commonly used in visual retrieval (e.g., ColQwen2) and significantly larger VLM for OCR-task. (2) We introduce and execute a rigorous experimental protocol that systematically evaluates both systems across precisely controlled levels of document degradation. By manually categorizing the documents based on their noise levels, we explicitly characterize each system's robustness and performance under conditions found in real-world deployment scenarios. (3) We propose and validate a novel semantic evaluation benchmark explicitly designed for assessing end-to-end performance in knowledge-intensive document question-answering tasks. Unlike conventional evaluations focused solely on retrieval accuracy, our benchmark incorporates automated semantic metrics including Exact Match, BLEU, ROUGE-1, and ROUGE-L.

2 RELATED WORK

This section gives an overview of various document retrieval works related to this research, highlighting the limitations and strengths of differing approaches across both OCR and VLM retrieval methods.

2.1 Vision-Language Models for Document Understanding

Early multimodal approaches to document understanding model both text and layout information. LayoutLM introduced unified transformer encoding textual tokens alongside spatial layout positions, greatly improving tasks such as form understanding and receipt parsing [18]. Further advancements incorporated visual features explicitly, such as DocFormer, which integrates text, layout, and visual embeddings within a transformer architecture [1]. More recently, LayoutLMv3 extended this methodology by pre-training both text and image modalities with a unified masking strategy, achieving state-of-the-art results across text-heavy and image-intensive document tasks [7]. Another strong advancement,

UDOP (Unifying Document Processing), introduced a generative vision-text-layout transformer model capable of different document AI tasks through a prompting-based framework, achieving state-of-the-art results across multiple document benchmarks [15]. While powerful, these methods still depend on the extracted textual inputs from OCR, supplemented with visual information and layout context.

2.2 OCR Integration with Retrieval-Augmented Generation

RAG systems rely on the quality of input documents for effective retrieval. Liu et al. (2023) showed that increases in OCR noise significantly reduces retrieval accuracy in various tested RAG applications. Zhang et al. (2024) introduced OHRBench, a benchmark for evaluating how various levels of OCR-induced noise affect RAG pipelines [21]. In that work, Zhang et al. found that even state-of-the-art OCR systems often fail to construct high-quality knowledge bases for retrieval when the original documents contain what they call semantic and formatting noise. Piryani et al. (2025) developed MultiOCR-QA, a multilingual dataset designed to test LLM performance on OCR-processed texts, in which the authors found that character-level OCR errors significantly reduce QA accuracy [12].

Transformer-based OCR models like Nougat and Llama 3.2 have improved the OCR bottleneck, but their effectiveness under real-world noise conditions is still limited. Understanding the interaction between OCR fidelity and retrieval outcomes is important for evaluating when traditional OCR pipelines are sufficient compared to newer approaches such as VLM-based retrieval (e.g., ColPali) that claim higher levels of robustness. This work situates itself within this space by directly comparing text embedding based RAG pipelines, driven by OCR, versus vision embedding RAG systems under varying document quality controls.

2.3 Vision-Language Models for Document Retrieval

VLMs offer an alternative to OCR methods by embedding entire document images, which includes text, layout, and visual context into a shared representation space. This approach bypasses the traditional OCR step, which could potentially reduce issues arising from text extraction error. CLIP [13] was one of the first large-scale VLMs to align visual and textual embeddings using contrastive learning. SigLIP [20], modified CLIP by introducing a sigmoid-based loss function that allows for more stable training and better scalability across noisy or individual text-image pairs. These models laid the groundwork for more document-specific architectures such as PaliGemma [2], which combines SigLIP’s image encoder with Gemma’s decoder to handle token-level document retrieval tasks.

ColBERT (Contextualized Late Interaction - BERT), expands upon BERT by adding a late interaction mechanism. The main difference from BERT is that this mechanism delays the interaction between query and document until the final scoring stage [8]. Similarly to BERT, ColBERT also stores embeddings at the word-token level for both query and document, instead of a single embedding at the document level, and then afterward scores similarity between each combination of word-document embedding. This in turn allows for a more in-depth similarity calculation between query and

document compared to its single-embedding per document counterpart, providing higher retrieval accuracy. ColPali [6] additionally builds on this approach by combining the PaliGemma model with ColBERT-style late interaction, to match visual document embeddings with token-level query embeddings.

2.4 ColPali and the ViDoRe Benchmark

ColPali is a RAG design that uses a VLM to bypass the traditional OCR-to-text pipeline by directly embedding image patches across the entire document [6]. By utilizing the late interaction mechanism seen in ColBERT, alongside PaliGemma’s ability to generate multi-vector embeddings for document images (patches) that capture both the textual and visual data inside each document, ColPali is able to efficiently match queries with relevant documents. Overall, ColPali aims to increase the accuracy of document retrieval, while also reducing training time by embedding text and visual chunks simultaneously in a single, end-to-end method. To evaluate performance, the authors of ColPali proposed the ViDoRe (Visual Document Retrieval) benchmark [6]. This dataset includes 127,346 high-quality PDF pages that generally all have well-structured layouts and legible content. Models are scored using metrics like NDCG@5, which tests the ability to retrieve the correct page given a specific query.

While ColPali demonstrates strong performance on this benchmark, the controlled nature of the datasets—high resolution, consistent formatting—limits its generalizability. Additionally, in its OCR-to-text RAG baseline, ColPali uses Tesseract, which is not considered state-of-the-art. This choice of a weaker baseline may have skewed the comparison in favor of the VLM approach. Our work addresses this gap by employing a state-of-the-art OCR model for a more robust and equitable comparison. This work also introduces a semantic evaluation step that compares the retrieved content’s answer to a ground truth answer, allowing the ability to move beyond exact page-match metrics and assess downstream utility in RAG pipelines.

2.5 Comparative Benchmarks for Vision-Language Models

Recent studies have introduced benchmarks to evaluate the strengths and weaknesses of VLMs across various document understanding tasks. HELM (Holistic Evaluation of Language Models) is a framework developed to provide a standardized, multidimensional evaluation of language models [11]. VHELM extends this framework to the vision-language setting [9]. LVLM-eHub provides a large-scale benchmark for evaluating instruction-tuned VLMs on multimodal tasks, to obtain a better understanding of generalization and issues with overfitting [17]. Other evaluations additionally explore architectural comparisons and reliability. For example, Mamba-based models have shown better performance compared to traditional transformer-based VLMs in vision-language tasks [16]. DeCC introduces a task decomposition framework to evaluate answer consistency and the robustness of VLMs [19]. Overall, these benchmarks highlight both the promise and limitations of VLMs in document-level reasoning, which is one of the primary focuses of this work on image-based retrieval. In addition, this work complements these efforts by evaluating VLM-based retrieval approaches, specifically in

the context of noisy, real-world documents and directly comparing them against OCR-based RAG pipelines.




Figure 1: Overview of the experimental pipeline comparing VLM-based and OCR-based RAG systems.

3 METHODOLOGY

In this section, we present the detailed methods and formal mathematical descriptions for the RAG pipelines compared in our study. We outline the VLM-based pipeline, the OCR-based pipeline, and the rigorous evaluation framework used to assess both retrieval and semantic accuracy.

3.1 System Architectures

We rigorously evaluate two distinct RAG system paradigms: a VLM-based approach that directly embeds images, and an OCR-based pipeline that extracts text before embedding.

3.1.1 VLM-Based RAG Pipeline. Our vision-based RAG pipeline follows the ColPali architecture proposed by Faysse et al. [6], specifically designed to bypass OCR preprocessing. The core architecture utilizes the PaliGemma model, integrating SigLIP-So400m as the image encoder and Gemma-2B as the textual encoder-decoder. Formally, a given document image D_i is partitioned into m non-overlapping patches $\{p_{i,1}, p_{i,2}, \dots, p_{i,m}\}$. Each patch $p_{i,j}$ is encoded by the VLM encoder f_{VLM} into embedding vectors:

$$\mathbf{e}_{i,j}^{(\text{VLM})} = f_{\text{VLM}}(p_{i,j}; \theta_{\text{VLM}}), \quad \mathbf{e}_{i,j}^{(\text{VLM})} \in \mathbb{R}^d,$$

where d denotes embedding dimensionality and θ_{VLM} the model parameters. The resulting multi-vector embedding representation is:

$$\mathbf{E}_i^{(\text{VLM})} = \{\mathbf{e}_{i,1}^{(\text{VLM})}, \dots, \mathbf{e}_{i,m}^{(\text{VLM})}\}.$$

At retrieval time, queries q are encoded using a textual embedding model g parameterized by ϕ :

$$\mathbf{q} = g(q; \phi), \quad \mathbf{q} \in \mathbb{R}^d.$$

Retrieval scoring employs ColBERT-style late-interaction, computing similarity scores between query and document patch embeddings:

$$s_i^{(\text{VLM})}(q, D_i) = \max_{j \in [1, m]} \frac{\mathbf{q}^\top \mathbf{e}_{i,j}^{(\text{VLM})}}{\|\mathbf{q}\| \|\mathbf{e}_{i,j}^{(\text{VLM})}\|}.$$

At the time of our experiments, the SOTA model from the ColPali family was ColQwen2 (7B), which we employed for the experiments presented here.

3.1.2 OCR-Based RAG Pipeline. Our OCR-based pipeline leverages advanced OCR extraction combined with semantic-rich text embeddings. Document images D_i first undergo OCR via an OCR model h_{OCR} :

$$T_i = h_{\text{OCR}}(D_i),$$

where T_i is the extracted textual content. This text is subsequently embedded using the Qwen embedding model $g(\cdot; \phi)$ to produce a single, dense semantic embedding:

$$\mathbf{e}_i^{(\text{OCR})} = g(T_i; \phi), \quad \mathbf{e}_i^{(\text{OCR})} \in \mathbb{R}^d.$$

Documents are stored in a vector database based on these embeddings. Query embedding follows the same procedure, ensuring consistent semantic representation. Retrieval scores for OCR embeddings use cosine similarity:

$$s_i^{(\text{OCR})}(q, D_i) = \frac{\mathbf{q}^\top \mathbf{e}_i^{(\text{OCR})}}{\|\mathbf{q}\| \|\mathbf{e}_i^{(\text{OCR})}\|}.$$

This pipeline's accuracy directly depends on OCR quality, making it potentially sensitive to visual degradation and complex layouts.

3.2 DocDeg: A Dataset for Degradation-Aware RAG Evaluation

We introduce **DocDeg**, a curated dataset specifically designed to evaluate Retrieval-Augmented Generation (RAG) systems under realistic visual degradation conditions. DocDeg comprises 4,196 diverse documents obtained from the U.S. Department of Energy's Office of Scientific and Technical Information (OSTI) and includes a wide variety of document types, such as academic papers, technical reports, memos, presentation slides, and handwritten notes, spanning both scanned and digitally generated quality based on OSTI source¹. While specific scanner settings such as resolution or color mode were not available in the source metadata, the breadth of the collection inherently captures a wide range of real-world conditions. The scanned documents often exhibit noise, artifacts, and quality loss frequently encountered in practical RAG deployments.

Each document page was manually labeled by two expert annotators into four distinct degradation categories reflecting visual clarity: Level 0 (native digital documents with no degradation), Level 1 (high-quality scans with minimal visual noise or artifacts), Level 2 (imperfect scans exhibiting moderate blur, faint text, or minor artifacts), and Level 3 (severely degraded pages featuring significant distortion, extensive noise, and considerable content loss).

¹<https://www.osti.gov/>

3.2.1 Feature Annotation. To enable a more detailed understanding of retrieval challenges, each document in the DocDeg dataset was manually annotated among a range of 12 structural and content features. Initially, 5,800 document pages were labeled by two annotators. After filtering out approximately 1,600 documents that lacked sufficient textual or visual information (e.g., blank pages, cover sheets, low-content slides), the final DocDeg dataset consists of 4,196 information-rich documents used in our experiments. Our analysis focuses on the degree of visual distortion or degradation. We judged quantified degradation on a scale of 0 to 3, with the following categories:

- **Level 0 (Digital Copy):** Native digital document with no visible degradation.
- **Level 1 (Perfect Scan):** High-quality scanned document with minimal or no noise/artifacts.
- **Level 2 (Imperfect Scan):** Noticeable visual defects such as mild blurring, faded or missing text, or scanning artifacts.
- **Level 3 (Severely Degraded):** Document is almost illegible, with heavy noise, distortion, or severe content loss.

A zoomed in example of each degradation level is provided in Figure 2.

3.2.2 Q&A Pair Generation and Benchmark Rigor. We constructed a rigorous evaluation benchmark by generating ten unique, nuanced question-answer pairs per document using Llama 3.3 (70B). We leveraged Sambanova Systems’ API calls to deploy the LLaMA models for Q&A metadata synthesis and retrieval evaluation. These pairs were carefully designed to capture specific factual details from each document, explicitly avoiding direct verbatim copying from the source text.

Using a LLM to generate a benchmark offers significant advantages in scalability and the creation of diverse, non-trivial queries that test deep comprehension. However, we acknowledge the potential risks, such as the introduction of factual errors (hallucinations) or ambiguous phrasing. To mitigate this, a subset of the generated Q&A pairs was manually validated for factual accuracy and clarity, confirming that the model produced contextually rich and relevant questions. Furthermore, as all systems in our study are evaluated against the same benchmark, any systemic noise or ambiguity is applied uniformly, thus preserving the fairness and validity of our comparative findings.

Leveraging these generated question-answer pairs, we performed two distinct evaluations of RAG system performance: (1) Document Retrieval Assessment, where each generated question was treated as a query to assess the system’s ability to correctly retrieve the corresponding source document, and (2) Semantic Answer Generation Evaluation, where we evaluated whether the RAG-augmented language model could accurately produce the reference answer provided in the generated pairs when conditioned on retrieved document context. This dual-evaluation setup enabled a comprehensive assessment of both retrieval precision and downstream semantic accuracy of the RAG systems under investigation.

3.3 Evaluation Framework

To rigorously assess both retrieval effectiveness and downstream semantic accuracy, we employ a dual evaluation framework consisting of standard retrieval metrics and a novel semantic answer evaluation benchmark.

3.3.1 Retrieval Evaluation Metrics. We utilize well-established retrieval metrics to quantitatively measure retrieval performance:

- **Mean Reciprocal Rank (MRR)** evaluates retrieval quality by averaging the reciprocal rank positions of correct documents: $MRR = \frac{1}{|Q|} \sum_{q \in Q} \frac{1}{r_q}$.
- **Recall@k** measures the proportion of relevant documents retrieved among the top-k results:

$$\text{Recall@k} = \frac{|\{\text{relevant documents in top } k\}|}{|\{\text{relevant documents total}\}|}.$$

- **Normalized Discounted Cumulative Gain (NDCG@k)** quantifies retrieval ranking quality by considering the position and relevance of documents retrieved:

$$\text{NDCG@k} = \frac{\sum_{i=1}^k \frac{2^{rel_i} - 1}{\log_2(i+1)}}{\sum_{i=1}^k \frac{2^{rel_i^{ideal}} - 1}{\log_2(i+1)}}.$$

The degradation labels were used to segment the dataset for all retrieval evaluations reported in this paper. Figure 2 provides a zoomed in example for each of the four human evaluated text degradation levels.

3.3.2 Semantic Answer Evaluation. In addition to retrieval metrics, we evaluate semantic accuracy using a robust end-to-end evaluation framework involving automated metrics for semantic comparison. Our semantic answer evaluation compares the answers generated by each RAG pipeline with those originally generated by Llama 3.3 (70B). We leveraged metrics such as Exact Match (EM), BLEU Score and ROUGE Scores (ROUGE-1, ROUGE-L) to evaluate the semantic accuracy.

3.4 Experimental Setup

RAG practitioners typically deploy systems “out of the box”, relying on pre-trained models without additional fine-tuning on their specific document collections due to practical constraints such as cost, time, expertise and data availability. With that in mind, our primary evaluation focuses on how each pipeline performs under realistic, visually degraded conditions using pre-trained models without performing task-specific fine-tuning on the DocDeg dataset for either the VLM-based or the OCR-based retrieval components. This approach allows us to assess the generalizability and robustness of these architectures when faced with unfamiliar document types and degradation patterns, which is crucial for understanding their practical utility. We report retrieval metrics across the four degradation levels defined in Section 3.2.1.

We report retrieval metrics—Mean Reciprocal Rank (MRR), Recall@5, and Normalized Discounted Cumulative Gain at rank 5 (NDCG@5)—across four levels of increasing document degradation on this dataset. In addition to DocDeg, we also conducted supplementary experiments on the ViDoRe benchmark (specifically, the DocVQA subset). This evaluation compared three approaches:

ry, MS I-7 Engineered Materials, Materials Science and Techno

dimensional correlation spectroscopy has improved on traditional samples to a variable amount they have been probed spectroscopically. These perturbed signals are at wavenumber axis and after the spectrum is created that the dynamic changes in the material are core of material degradation can that Noda proposed.² This technique studies because of improved of complex or overlapping traditional analytical method utilized since its inception, and lies, as well as even wine chemistries of band separation. Two-tion analysis has also been demonstrated with different spectroscopic technique has increasing avenues used in many other fields in the

Generalized Two-Dimensional (2D-COS). This technique is spectra and overlapped bands in traditional IR spectroscopy, the correlation in bands between methods (IR, Raman, etc.). A unique feature is that it can be univocally interpreted by any of the previous external perturbations before the measurement (Figure 1). By increasing the perturbation, distinct changes are observed by overlaying multiple spectra. Changes in a material. Specifically, generate a 2D correlation correlation dynamic spectrum.²

Figure 1 illustrates the general process to obtain a sample already affected by the external field. It shows an 'Electro-Magnetic Probe (IR, UV, etc.)' with an arrow pointing to a box labeled 'S'. Above the box is the word 'External'.

rior Transform Infrared was created by Isao Noda, shows an increase in resolution spectra against a sec-axis.¹ The desire for this technical spectroscopy became R field. As 2D spectroscopy ical spectroscopy began to be 2D-FTIR arose from relaxa-

ing experience in the context of existing behavior. One of the key features of ATHEANA is that *error-forcing contexts*¹ (equipment/material conditions and performance shaping) is under which *unsafe actions* (UAs) can occur. (See the Manual [Taylor et. al., 1997] and Implementation Guide for more information.)

A relies heavily on the analysis of operational events (A) in order to generate creative thinking about possible EFCs, a process called the Human Events Reference for ATHEANA (HERA) methodology. This report documents the initial development

ment effort is a follow-on activity to two earlier db in Action Classification Scheme (HACS) (Barriere, 4) pre-data ATHEANA. It was patterned on existing titation Process (HPPI) (Paradies, Unger, Haas, and zed Library for Assessing Reactor Reliability (NUCL), so represented the commonly accepted aspects of h phases, and error types. This strong link between l inventions of HRA/PRA, however, rendered the db in IRA framework of ATHEANA. Thus, a second evol system Event Classification Scheme or HSECS, was u ATHEANA. (See Cooper, Luckas, and Wreathall,


at a definition for the highlighted term is available in the glossar

ever form. The upper limit of 2000 km is to be recognizable as a granule, the distance less than the distance from optical features are integral scale features of the turbulence; interesting that the thermal microscale diameters, not much smaller than the convective velocity microscale is one centimeter smaller.

steady-state models have been notable for ignoring the physics of the granulation bar and quasi-periodic. Discovery of the role of individual granules was crucial for multidimensional ensemble-averaged models. (These models are also suspect because they average the buoyant rise or other time-dependent processes during important physics.

ization is that the granulation is strictly proportional to the pressure scale height. Therefore its use to calibrate phenomenological models can be justified. In particular, the fact that one pressure scale height in diameters corresponds to the commonly-used so-called "granulation length scale" can be no larger than approximately one solar radius. A recent calculation suggests any connection between granulation and pressure scale height.

ation suggested by this model requires additional beliefs about the supergranulation: if granules are thermally-driven convection cells (frequently speculated to be He ionization fronts), the granulation behavior is a giant bubble rising from below, and the supergranulation is a giant bubble rising from below.



important parameters affecting FRC formation in FRC: the bias field, and the timing of the ringing theta-pi-pi was observed by using an axial array of side-viewing traction strength, an end-viewing visible (500-600 nm) to monitor azimuthal symmetry during the preionization, and the usual exclude-1 flux diagnostic to monitor the quality of formation. Recently, an end-view with 2 μ s exposure time has been used (in collaboration with symmetry of the $T_e \sim 20 \sim 30$ eV plasma after which good confinement could be obtained, measure the density distribution indicate that the density profile is visible framing photographs show a current sheet wide



Figure 2: Zoomed-in view of documents with different degradation levels

- **ColQwen (baseline):** ColPali retrieval with their SOTA model, ColQwen2-v1, fine-tuned on the ViDoRe dataset.
- **OCR+Qwen2 (Nougat):** OCR extraction by rerunning the ViDoRe images through the Nougat OCR model, followed by Qwen2 7B Instruct without fine-tuning.
- **OCR+Qwen2 (Llama):** OCR extraction by rerunning the ViDoRe images through the Llama3.2 90B OCR model, followed by Qwen2 7B Instruct, again without fine-tuning.

4 RESULTS AND DISCUSSION

In this section, the findings from our experiments are presented, an analysis of performance differences is given, and the resulting implications are discussed in the context of document understanding and retrieval systems.

4.1 Retrieval Performance Across Document Quality Levels

To evaluate the robustness and accuracy of each pipeline across realistic conditions, we systematically assessed their retrieval performance over the four incremental document degradation levels mentioned in Section 3.2.1. We report Mean Reciprocal Rank (MRR), Recall@5, and Normalized Discounted Cumulative Gain at rank 5 (NDCG@5).

Table 1 compares a vision-language model (VLM)-based pipeline against two OCR-based pipelines, one using nougat OCR and another using Llama 3.2 for OCR.

As illustrated in Table 1 and Figure 3, the Llama OCR-Based pipeline consistently outperformed the VLM-Based pipeline on retrieval accuracy across all degradation levels evaluated on the DocDeg dataset for all metrics. Both systems show declines in accuracy as document quality decreases, though the OCR-Based

Table 1: Retrieval performance metrics across document degradation levels.

VLM-Based = Vision-Language Model pipeline, **OCR-Based (Nougat)** = OCR pipeline with Nougat OCR, **OCR-Based (Llama)** = OCR pipeline with Llama 3.2 OCR.

Distortion Level	Questions	Metric	VLM-Based	OCR-Based (Nougat)	OCR-Based (Llama)
Level 0	11190	MRR	0.2971	0.2327	0.5151
		Recall@5	0.4057	0.3013	0.6415
		NDCG@5	0.3242	0.2721	0.5468
Level 1	19154	MRR	0.2440	0.3218	0.4857
		Recall@5	0.3573	0.4298	0.6525
		NDCG@5	0.2722	0.3627	0.5275
Level 2	6864	MRR	0.2012	0.3119	0.4579
		Recall@5	0.3003	0.4093	0.6104
		NDCG@5	0.2259	0.3502	0.4961
Level 3	4748	MRR	0.2098	0.2295	0.4520
		Recall@5	0.3112	0.3010	0.5925
		NDCG@5	0.2350	0.2612	0.4872
Total	41956	MRR	0.2471	0.2858	0.4852
		Recall@5	0.3554	0.3774	0.6359
		NDCG@5	0.2740	0.3212	0.5229

scores still significantly outperformed the VLM-Based ones even under severe degradation.

Comparing Nougat to Llama further highlights that the specific OCR engine used within the OCR-Based pipeline has a substantial impact on overall retrieval performance. The Llama 3.2 OCR engine consistently achieved higher retrieval metrics than Nougat across all distortion levels. This underscores the critical role that OCR quality plays in downstream retrieval tasks, particularly when deploying OCR-to-RAG systems in realistic, noisy document environments.




Figure 3: Retrieval performance across document quality levels

Table 2: Retrieval performance metrics across document degradation levels (slides removed).

VLM-Based = Vision-Language Model pipeline, **OCR-Based (Nougat)** = OCR pipeline with Nougat OCR, **OCR-Based (Llama)** = OCR pipeline with Llama 3.2 OCR.

Distortion Level	Questions	Metric	VLM-Based	OCR-Based (Nougat)	OCR-Based (Llama)
Level 0	3784	MRR	0.1905	0.2819	0.3769
		Recall@5	0.2826	0.3979	0.5322
		NDCG@5	0.2134	0.3109	0.4157
Level 1	16974	MRR	0.2347	0.3416	0.4782
		Recall@5	0.3476	0.4593	0.6502
		NDCG@5	0.2628	0.3711	0.5213
Level 2	6794	MRR	0.2005	0.3139	0.4557
		Recall@5	0.2997	0.4122	0.6086
		NDCG@5	0.2252	0.3386	0.4941
Level 3	4748	MRR	0.2098	0.2295	0.4520
		Recall@5	0.3112	0.3010	0.5925
		NDCG@5	0.2997	0.2474	0.4872
Total	32300	MRR	0.2186	0.3123	0.4578
		Recall@5	0.3244	0.4189	0.6192
		NDCG@5	0.2449	0.3390	0.4982

Nougat outperformed the VLM-based RAG pipeline on levels one through three and on the weighted average. Interestingly, Nougat performed worse on the highest quality documents. 1 shows retrieval results with slideshows removed from the dataset. Both Llama and ColQwen saw a reduction in retrieval accuracy on level zero documents when slides were excluded. Meanwhile, Nougat’s performance on the slideshow-free data improved.

Although ColPali retrieves the top-5 visually relevant documents based on query-image similarity, it does not natively support answer generation. To use ColPali in a full RAG pipeline, an additional step is required: either (1) apply OCR to the retrieved document images and pass the extracted text into a language model, or (2) use a vision-language model capable of question answering over images. Both options introduce additional computational cost, either in the form of OCR latency or VLM inference time. For this study’s semantic answer evaluation, we used the OCR text from the images retrieved by ColPali as context for the LLM, in order to enable a fair comparison across pipelines.

4.2 Semantic Answer Accuracy

In addition to retrieval accuracy, we provide a comprehensive evaluation of the pipelines’ end-to-end performance in question-answering scenarios using Exact Match, BLEU, ROUGE-1, and ROUGE-L (Table 3). In addition to automated metrics, a small-scale manual review of approximately 40 generated question-answer pairs was conducted for qualitative verification. While not exhaustive, this review confirmed the quality and relevance of the question-answer pairs to the documents.

Table 3: Semantic answer quality across document degradation levels.

VLM-Based = ColQwen2, **OCR-Based** = Llama 3.2 90B for OCR + Qwen2.

Distortion Level	Questions	Metric	VLM-Based	OCR-Based
Level 0	11190	Exact Match	0.0046	0.0080
		Average BLEU	0.0589	0.0753
		Average ROUGE-1	0.2088	0.2547
		Average ROUGE-L	0.1912	0.2370
Level 1	19154	Exact Match	0.0038	0.0066
		Average BLEU	0.0628	0.0804
		Average ROUGE-1	0.2167	0.2643
		Average ROUGE-L	0.1960	0.2430
Level 2	6864	Exact Match	0.0033	0.0057
		Average BLEU	0.0587	0.0751
		Average ROUGE-1	0.1975	0.2410
		Average ROUGE-L	0.1785	0.2213
Level 3	4748	Exact Match	0.0038	0.0066
		Average BLEU	0.0595	0.0762
		Average ROUGE-1	0.2044	0.2493
		Average ROUGE-L	0.1826	0.2263
Total	41956	Exact Match	0.0039	0.0068
		Average BLEU	0.0607	0.0777
		Average ROUGE-1	0.2100	0.2556
		Average ROUGE-L	0.1903	0.2355

Results from Table 3 show the semantic answer evaluation. Consistent with the results obtained from retrieval, the OCR-Based pipeline obtained higher scores when compared to the VLM-Based pipeline on every tested metric and degradation level. A notable difference between these results and the retrieval results, is that there does not seem to be an obvious decrease in score as degradation

level increases. Overall, the results suggest that the potential benefits of the VLM's ability to handle visual nuance did not translate to improved semantic performance when compared to the superior results from the OCR-Based pipeline.

4.3 Computational Efficiency Analysis

A comparative analysis of indexing time, retrieval latency, and memory consumption for both pipelines is summarized in Table 4. The results reveal that the VLM-based pipeline is more efficient in both embedding generation and retrieval latency compared to the OCR-based pipeline. The VLM-based system achieved lower embedding times (0.4739s per document compared to 0.5503s) and substantially reduced retrieval latency (0.0010s per query vs. 0.0311s per query). Additionally, memory consumption per 1,000 documents was notably lower for the VLM-based pipeline (1.38 GB) relative to the OCR-based pipeline (9.5 GB). This substantial memory efficiency can be primarily attributed to the avoidance of storing extensive OCR-derived textual embeddings.

All experiments not involving Llama-based models were conducted on a node equipped with two NVIDIA A100 GPUs (40GB VRAM each), an AMD EPYC 7502 CPU, and 251 GB of system memory. However, the SentenceTransformer model was executed on a single A100 throughout. The second GPU remained unused. For OCR extraction and question-answer pair generation involving Llama 3.2 and Llama 3.3 models, we utilized a SambaNova SN40L platform with integrated RDU accelerators. Retrieval latency, embedding times, and memory usage metrics reported in Table 4 reflect performance measured on these respective systems. Embedding of text and images was done with a batch size of 1.

Table 4: Computational Efficiency Comparison

Metric	VLM-Based	OCR-Based
OCR Time (per document)	N/A	6s
Embedding Time (per document)	.4739s	.5503s
Retrieval Latency (per query)	0.04252s	.0311s
Memory Usage (per 1k documents)	1.38 GB	9.5 GB

4.4 Overall assessment on DocDeg Dataset

We leverage a RADAR plot to showcase overall performance of VLM vs OCR RAG shown in Figure 4. Each axis is a linearly rescaled version of a concrete, published metric—chosen so that a score of 10 corresponds to the best value observed across *all* pipelines and degradation levels, and a score of 0 would correspond to a value at least two standard deviations worse than the worst observed.² Specifically:

- (1) **Clean-document retrieval** = $10 \times \frac{MRR_{d=0} - \mu_{MRR}}{\sigma_{MRR}}$, clipped to $[0, 10]$.
- (2) **Noisy-document retrieval** = $10 \times \frac{MRR_{d=3} - \mu_{MRR}}{\sigma_{MRR}}$.
- (3) **Semantic answer quality** = $10 \times \frac{ROUGE-L_{total} - \mu_R}{\sigma_R}$, where μ_R, σ_R are the mean and s.d. of ROUGE-L across all runs.

²Because no pipeline actually reached that extreme, the lowest displayed score is ≈ 5.5 .




Figure 4: Comparative Analysis of RAG System Capabilities

- (4) **Processing speed** = $10 \times \frac{\min(T_{ret})}{T_{ret}}$, where T_{ret} is per-query latency; lower is better.
- (5) **Memory efficiency** = $10 \times \frac{\min(M)}{M}$, where M is memory per 1k documents.

All raw values (MRR , $ROUGE-L$, latency, memory) are reported in Tables 1 and 4. The resulting normalised scores for VLM, OCR+Nougat, and OCR+Llama are plotted without additional smoothing. This transparent mapping ensures that the radar visualisation is a faithful, monotonic transformation of the quantitative results.

The radar plot provides a comprehensive visualization of the relative strengths and weaknesses between VLM-based and OCR-based RAG systems across five critical dimensions. The VLM-based approach demonstrates exceptional performance in processing speed (9.5/10) and memory efficiency (9.0/10), reflecting its significantly faster retrieval latency (0.001s vs 0.031s) and lower memory footprint (1.38GB vs 9.5GB per 1k documents). However, it shows limitations in both clean document retrieval (6.5/10) and noisy document retrieval (6.0/10). In contrast, the OCR-based pipeline with Llama 3.2 exhibits superior performance in clean document retrieval (9.0/10), noisy document retrieval (8.5/10), and semantic answer quality (8.5/10), at the cost of greater computational demands. The Nougat OCR-based system offers a middle ground across most metrics. These findings reveal a clear trade-off: while modern OCR-based pipelines leveraging advanced models like Llama 3.2 provide better retrieval accuracy and semantic quality in realistic document settings, VLM-based approaches offer significant advantages in computational efficiency and resource utilization.

4.5 ViDoRe Evaluation Results

To further evaluate retrieval performance under cleaner document conditions, we conducted supplementary experiments on the DocVQA subset of the ViDoRe benchmark. We compared three retrieval pipelines: (1) the ColQwen2 model, representing the ViDoRe authors' best fine-tuned vision-language system based on ColPali; (2) an OCR-based pipeline using Alibaba-NLP/gte-Qwen2-7B-instruct with Nougat OCR outputs; and (3) the same OCR-based pipeline but preprocessing the images into text with the more advanced Llama 3.2 90B vision model.

Table 5 summarizes the key retrieval metrics across the three approaches.

Table 5: ViDoRe DocVQA Retrieval Comparison

Metric (Top-5)	ColQwen2	OCR+Qwen2 (Nougat)	OCR+Qwen2 (Llama)
NDCG@5	0.6027	0.2930	0.3774
MAP@5	0.5813	0.2762	0.3484
MRR@5	0.5851	0.2786	0.3484
Recall@5	0.6674	0.3437	0.4656

ColQwen2, which was fine-tuned on DocVQA, achieves the strongest retrieval performance, with an NDCG@5 of 0.6027 and a Recall@5 of 0.6674.

While a few ViDoRe pages have some minor noise, most pages in ViDoRe are clean and contain little text. The benchmark remains challenging not because of OCR quality, but because its queries are only loosely connected to the page text—typically short key phrases or vague prompts (e.g., “what was the dividend payout in 2012?”) rather than full natural-language questions. At the same time, many ViDoRe pages contain sparse fragments of text rather than dense paragraphs, which makes it harder to establish semantic links between queries and documents.

ColQwen2 succeeds because it has been fine-tuned on ViDoRe’s own query–page pairs, so its embeddings have learned to capture precisely these idiosyncratic connections. In contrast, our Nougat and Llama OCR pipelines rely on a general-purpose text encoder that has never seen ViDoRe queries; they fall behind primarily because their embeddings are semantically misaligned, not because their OCR output is faulty. Indeed, Nougat and Llama 3.2 produce similarly poor but not identical retrieval scores, confirming that on relatively clean PDFs with sparse text OCR quality is no longer the bottleneck, semantic alignment between the query and the page embeddings is.

4.6 OCR Output Divergence Analysis

Table 6: Average Levenshtein Distance Between OCR Engine Outputs for documents for each level. A lower score indicates higher similarity between the extracted text from the same source document.

Level	L ↔ N	L ↔ T	T ↔ N
Level 0	734.63	463.66	633.26
Level 1	747.27	395.18	621.60
Level 2	1478.97	900.51	1085.02
Level 3	2609.10	1940.70	1490.02

Average Levenshtein distance between OCR outputs (lower = more similar).
L: Llama-3.2 90B, N: Nougat, T: Tesseract

To better isolate the impact of the OCR engine itself, separate from the downstream retrieval and generation tasks, we conducted a direct comparison of the textual outputs from our primary OCR




Figure 5: OCR Output Divergence by Document Degradation Level. The chart shows the average Levenshtein distance between pairs of OCR models. Higher bars indicate greater divergence in the extracted text. The divergence escalates dramatically for severely degraded (Level 3) documents.

models, Llama 3.2 and Nougat, as well as Tesseract, which was used in the OCR-based RAG baseline reported in the ColPali paper. We used the Levenshtein distance to quantify the difference, or divergence, between the text extracted by each pair of OCR models for the same document. A higher score signifies greater divergence (the texts are more different), while a lower score indicates the models are in closer agreement.

This analysis, detailed in Table 6 and visualized in Figure 5, reveals several critical findings about the nature of OCR “translation.”

At low degradation levels (Level 0 and 1), the outputs of Llama 3.2 and Tesseract are the most aligned, exhibiting the lowest Levenshtein distance. This suggests that on cleaner documents, the advanced capabilities of Llama and the established algorithms of Tesseract converge on a similar, accurate textual representation. Nougat, in contrast, produces the most distinct output in these scenarios, possibly explained by the fact that nougat produces MultiMarkdown (.mmd) format rather than pure text. As noise increases (Level 2 and 3), all models begin to diverge more significantly.

4.6.1 Implications for RAG. This analysis provides a clear, quantitative explanation for the performance differences observed in our end-to-end RAG evaluations. The “translation” from image to text is not a uniform process. The significant divergence, especially on noisy documents, means that the knowledge base being fed into the retrieval system can vary substantially based solely on the OCR engine used. An error, hallucination, or omission in one OCR output may not exist in another, directly impacting which documents are retrieved and what information is available for answer generation. This underscores that the selection of an OCR engine is not merely a preprocessing step but a critical decision that profoundly shapes the quality and reliability of the entire RAG pipeline, particularly when dealing with real-world, imperfect documents.

5 CONCLUSION AND FUTURE WORK

Our study rigorously compared OCR-based and VLM-based RAG pipelines, evaluating their performance and robustness across real-world degraded document scenarios. Our primary contribution is not simply to confirm that high-quality OCR benefits RAG, but to quantify the critical trade-offs between these two competing

paradigms. Our results demonstrate that while VLM-based retrieval offers significant advantages in computational efficiency and memory usage, state-of-the-art OCR-based pipelines provide a substantial and consistent lead in retrieval accuracy and downstream semantic quality, even on severely degraded documents.

Specifically, we found that OCR-based pipelines, despite incurring an upfront OCR cost at ingestion, are faster at query time than the evaluated VLM-based late-interaction architecture. Furthermore, the VLM approach does not inherently support end-to-end question answering; it requires an additional OCR or VQA step to generate a final answer, adding further latency. These findings highlight a clear trade-off: practitioners must weigh the superior accuracy and semantic fidelity of modern OCR-RAG systems against the lower resource footprint of VLM-based alternatives.

For future work, several avenues warrant exploration. To address the generalizability of our findings, the evaluation could be extended to other VLM retrieval architectures beyond ColPali and expanded to include multilingual document collections. Finally, while our experiments used pre-trained models to reflect common deployment scenarios, a valuable next step would be to compare fine-tuned OCR-based and VLM-based systems to determine the performance ceiling of each approach.

6 ACKNOWLEDGEMENTS

This manuscript has been approved for unlimited release and has been assigned LA-UR-25-24280. This research was funded by the LANL ASC grant AI4Coding and the LANL Institutional Computing Program, supported by the U.S. DOE NNSA under Contract No. 89233218CNA000001. We also thank Sambanova Systems for providing access to API calls for LLM inference utilized in this work

REFERENCES

- [1] Srikanth Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R. Manmatha. 2021. DocFormer: End-to-End Transformer for Document Understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 973–983. arXiv:2106.11539 [cs.CV]
- [2] Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschanen, Emanuele Bugliarello, Thomas Unterthiner, Daniel Keysers, Skanda Kopula, Fangyu Liu, Adam Grycner, Alexey Gritsenko, Neil Houlsby, Manoj Kumar, Keran Rong, Julian Eisenschlos, Rishabh Kabra, Matthias Bauer, Matko Bošnjak, Xi Chen, Matthias Minderer, Paul Voigtlaender, Ioana Bica, Ivana Balazevic, Joan Puigcerver, Pinelopi Papalampidi, Olivier Henaff, Xi Xiong, Radu Soricut, Jeremiah Harmsen, and Xiaohua Zhai. 2024. PaliGemma: A versatile 3B VLM for transfer. arXiv preprint. arXiv:2407.07726 [cs.CV]
- [3] Manish Bhattarai, Ryan Barron, Maksim E. Eren, Minh N. Vu, Vesselin Grantcharov, Ismael Ismael, Valentin Stanev, Cynthia Matuszek, Vladimir I Valtchinov, Kim Rasmussen, and Boian S. Alexandrov. 2025. HEAL: Hierarchical Embedding Alignment Loss for Improved Retrieval and Representation Learning. In *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing*. Association for Computational Linguistics, Albuquerque, New Mexico, USA, 205–214. <https://aclanthology.org/2025.knowledge-nlp-1.19/>
- [4] Manish Bhattarai, Javier E. Santos, Shawn Jones, Ayan Biswas, Boian Alexandrov, and Daniel O'Malley. 2024. Enhancing Code Translation in Language Models with Few-Shot Learning via Retrieval-Augmented Generation. In *2024 IEEE High Performance Extreme Computing Conference (HPEC)*. 1–8. <https://doi.org/10.1109/HPEC62836.2024.10938485>
- [5] Manish Bhattarai, Minh N. Vu, Javier E. Santos, Ismael Ismael, and Daniel O'Malley. 2025. Enhancing Cross-Language Code Translation via Task-Specific Embedding Alignment in Retrieval-Augmented Generation. In *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing*. Association for Computational Linguistics, Albuquerque, New Mexico, USA, 107–117. <https://aclanthology.org/2025.knowledge-nlp-1.8/>
- [6] Manuel Paysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, Céline Hudelot, and Pierre Colombo. 2024. ColPali: Efficient Document Retrieval with Vision Language Models. In *The Twelfth International Conference on Learning Representations (ICLR)*. arXiv:2402.01410 [cs.CV]
- [7] Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. LayoutLMv3: Pre-training for Document AI with Unified Text and Image Masking. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*. 4233–4242. arXiv:2204.08387 [cs.CV]
- [8] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient Passage Search via Contextualized Late Interaction over BERT. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 39–48. arXiv:2004.12832 [cs.IR]
- [9] Tony Lee, Haoqin Tu, Chi Heem Wong, Wenhao Zheng, Yiyang Zhou, Yifan Mai, Josselin Somerville Roberts, Michihiro Yasunaga, Huaxiu Yao, Cihang Xie, and Percy Liang. 2024. VHELM: A Holistic Evaluation of Vision Language Models. arXiv preprint. arXiv:2410.07112 [cs.CV]
- [10] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*. 9459–9474. arXiv:2005.11401 [cs.CL]
- [11] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Karan Santhanam, Jared Quincy Davis, Yura Alexandrov, Yacine Jernite, et al. 2022. Holistic Evaluation of Language Models. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022) Datasets and Benchmarks Track*.
- [12] Bhawna Piryani, Jamshid Mozafari, Abdelrahman Abdallah, Antoine Doucet, and Adam Jatowt. 2025. MultiOCR-QA: Dataset for Evaluating Robustness of LLMs in Question Answering on Multilingual OCR Texts. arXiv preprint. arXiv:2502.16781 [cs.CL]
- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 8748–8763. arXiv:2103.00020 [cs.CV]
- [14] Ray Smith. 2007. An Overview of the Tesseract OCR Engine. In *ICDAR '07: Proceedings of the Ninth International Conference on Document Analysis and Recognition*. IEEE Computer Society, Washington, DC, USA, 629–633. <https://storage.googleapis.com/pub-tools-public-publication-data/pdf/33418.pdf>
- [15] Zineng Tang, Ziyi Yang, Guoxin Wang, Yuwei Pang, Yang Liu, Chenguang Zhu, Michael Zeng, Cha Zhang, and Mohit Bansal. 2023. Unifying Vision, Text, and Layout for Universal Document Processing. In *CVPR*. 19254–19264.
- [16] Roger Waleffe, Wonmin Byeon, Duncan Riach, Brandon Norick, Vijay Korthikanti, Tri Dao, Albert Gu, Ali Hatamizadeh, Sudhakar Singh, Deepak Narayanan, Garvit Kulshreshtha, Vartika Singh, Jared Casper, Jan Kautz, Mohammad Shoeybi, and Bryan Catanzaro. 2024. An Empirical Study of Mamba-based Language Models. arXiv preprint. arXiv:2406.07887 [cs.CL]
- [17] Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Qiao, and Ping Luo. 2023. LVLm-eHub: A Comprehensive Evaluation Benchmark for Large Vision-Language Models. arXiv preprint. arXiv:2306.09265 [cs.CV]
- [18] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. LayoutLM: Pre-training of Text and Layout for Document Image Understanding. In *Proceedings of ACM SIGKDD*. 1192–1200.
- [19] Qian Yang, Weixiang Yan, and Aishwarya Agrawal. 2024. Decompose and Compare Consistency: Measuring VLMs' Answer Reliability via Task-Decomposition Consistency Comparison. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 3613–3627. arXiv:2404.16003 [cs.CV]
- [20] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. SigMoid Loss for Language Image Pre-Training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 11941–11952. <https://doi.org/10.1109/ICCV51070.2023.01100> arXiv:2303.15343 [cs.CV]
- [21] Junyuan Zhang, Qintong Zhang, Bin Wang, Linke Ouyang, Zichen Wen, Ying Li, Ka-Ho Chow, Conghui He, and Wentao Zhang. 2024. OCR Hinders RAG: Evaluating the Cascading Impact of OCR on Retrieval-Augmented Generation. arXiv preprint. arXiv:2412.02592 [cs.CL]