

IMRNNs: An Efficient Method for Interpretable Dense Retrieval via Embedding Modulation

Yash Saxena[†], Ankur Padia[†], Kalpa Gunaratna[‡], Manas Gaur[†]

[†]University of Maryland, Baltimore County, Baltimore, Maryland, USA

[‡]Independent Researcher, San Jose, California, USA

{ysaxena1, pankur1, manas}@umbc.edu, gunaratnak@acm.org

Abstract

Interpretability in black-box dense retrievers remains a central challenge in Retrieval-Augmented Generation (RAG). Understanding how queries and documents semantically interact is critical for diagnosing retrieval behavior and improving model design. However, existing dense retrievers rely on static embeddings for both queries and documents, which obscures this bidirectional relationship. Post-hoc approaches such as re-rankers are computationally expensive, add inference latency, and still fail to reveal the underlying semantic alignment. To address these limitations, we propose Interpretable Modular Retrieval Neural Networks (IMRNNs), a lightweight framework that augments any dense retriever with dynamic, bidirectional modulation at inference time. IMRNNs employs two independent adapters: one conditions document embeddings on the current query, while the other refines the query embedding using corpus-level feedback from initially retrieved documents. This iterative modulation process enables the model to adapt representations dynamically and expose interpretable semantic dependencies between queries and documents. Empirically, IMRNNs not only enhances interpretability but also improves retrieval effectiveness. Across the BEIR Benchmark, applying our method to standard dense retrievers yields average gains of +6.35% in nDCG, +7.14% in recall, and +7.04% in MRR over state-of-the-art baselines. These results demonstrate that incorporating interpretability-driven modulation can both explain and enhance retrieval in RAG systems.

1 Introduction

Retrieval-Augmented Generation (RAG) systems have emerged as a dominant paradigm for grounding large language models (LLMs) in factual, domain-specific knowledge (Lewis et al., 2020; Glass et al., 2022). At the heart of these systems

lies the *initial retriever*, responsible for selecting candidate documents from massive corpora before downstream re-ranking and generation. This component defines both the *efficiency* and *trustworthiness* of the entire pipeline. However, current dense retrieval methods face a fundamental limitation: they operate with *static embeddings* that encode queries and documents into fixed vector representations, preventing semantic adaptation between them at inference time.

This static nature creates two intertwined problems. First, it limits *retrieval performance*. When query and document embeddings cannot adapt to each other’s semantic context, the retriever struggles to capture context-sensitive relevance signals, particularly for complex or ambiguous queries (Li et al., 2021). Second, and more critically underexplored, it lacks *interpretability*, unlike lexical methods (e.g., BM25) that provide transparent term-matching explanations (Robertson and Zaragoza, 2009a), dense retrievers function as black boxes, preventing users from understanding which aspects of a query influenced retrieval decisions or how specific document features contributed to ranking outcomes (Zhou et al., 2024). This opacity is particularly problematic in high-stakes domains such as healthcare (Munnangi et al., 2025), legal research (Magesh et al., 2025), and finance (Kim et al., 2025), where understanding retrieval decisions is essential for trust and accountability.

Recent research has begun addressing retrieval performance through lightweight adaptation mechanisms. *Retrieval adapters*, such as SEARCH-ADAPTOR (Yoon et al., 2024a) and HYPEN-CODER (Killingback et al., 2025a) have applied learned transformations over frozen embeddings, while DIME (Campagnano et al., 2025a) leverages Matryoshka representations to select important embedding subspaces. However, these methods share

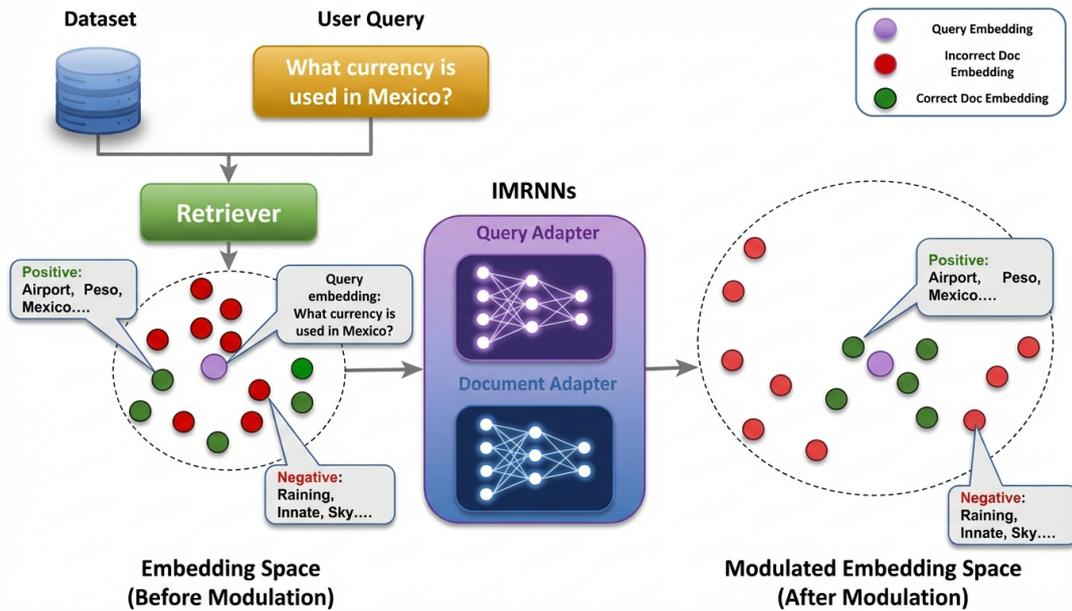


Figure 1: Example illustrating IMRNNs’s modulation mechanism. Starting from a static embedding space produced by the retriever, IMRNNs bidirectionally modulates query and document embeddings to form a modulated embedding space, drawing relevant documents closer to the query while pushing irrelevant ones away. The modulation is interpretable via modulation vectors and their associated key tokens: positive documents align with keywords such as Peso and Mexico, while negative documents align with Raining and Sky.

a critical gap: *they prioritize computational efficiency and marginal performance gains while providing no semantic-level interpretability.*

SEARCH-ADAPTOR learns dataset-level correction matrices for queries and documents, adjusting scores uniformly but obscuring which semantic aspects drive changes. HYPENCODER generates query-specific MLP scorers to replace cosine similarity, yet these scorer weights do not reveal which query concepts determine relevance. DIME selects top dimensions by ℓ_2 magnitude (e.g., 256/1024), showing which dimensions matter but not their semantic meaning. All existing adapters operate on dimensions or learned functions without translating to human-understandable concepts.

We introduce **Interpretable Modular Retrieval Neural Networks (IMRNNs)**, the first retrieval adapter providing simultaneous performance gains and multi-level interpretability without external explanation methods. IMRNNs comprises two lightweight MLPs on static embeddings: a *Query Adapter* modulating document embeddings \mathbf{d} conditioned on query \mathbf{q} (pulling relevant documents closer, pushing irrelevant ones away), and a *Document Adapter* aggregating corpus-level signals to adapt the query embedding. IMRNNs uses a bidirectional modulation mechanism to achieve three forms of interpretability *ignored* in prior work:

- *Structural interpretability*: The modulation mechanism uses explicit affine transformations $\mathbf{W}_q \mathbf{d} + \mathbf{b}_q$ and $\bar{\mathbf{W}}_d \mathbf{q} + \bar{\mathbf{b}}_d$, where \mathbf{W}_q and $\bar{\mathbf{W}}_d$ are learnable weight matrices and \mathbf{b}_q and $\bar{\mathbf{b}}_d$ are learned bias vectors. These parameters are directly observable. Unlike multi-layer neural scorers, users can inspect exactly which mathematical operations transformed the embeddings (Wang et al., 2023; Arendt et al., 2021).
- *Attribution-level interpretability*: By computing the difference $\Delta \mathbf{d} = \mathbf{d}_{\text{mod}} - \mathbf{d}_{\text{orig}}$ and $\Delta \mathbf{q} = \mathbf{q}_{\text{mod}} - \mathbf{q}_{\text{orig}}$, where \mathbf{d}_{mod} and \mathbf{q}_{mod} are the modulated embeddings, we obtain the exact change vector induced by modulation. This reveals precisely which embedding dimensions increased or decreased, and by how much, enabling dimension-level attribution of retrieval decisions (Zhou et al., 2022; Zhang et al., 2022; Calderon and Reichart, 2025).
- *Semantic-level interpretability*: We back-project the change vectors $\Delta \mathbf{d}$ and $\Delta \mathbf{q}$ from the adapter’s working space to the original encoder’s token embedding space using the Moore-Penrose pseudoinverse (Barata and Hussein, 2012). Let P^+ denote this pseudoinverse. By computing cosine similarity between the back-projected vector and every token embedding, we identify tokens whose semantics align with the modulation direc-

tion, revealing *what semantic concepts* drove the retrieval decision (Rajagopal et al., 2021; Sajjad et al., 2022).

Crucially, the identified tokens are not post-hoc explanations but directly correspond to the mathematical transformations that changed rankings. This contrasts with prior interpretability approaches (Yuksel and Kamps, 2025a; Llordes et al., 2023a) that analyze model behavior after the fact rather than exposing the mechanism itself. IMRNNs adapts per-query while preserving cosine similarity’s efficiency. Figure 1 shows the complete end-to-end workflow.

Contributions. Our work makes the following contributions:

- We propose IMRNNs, the **first retrieval adapter** with structural, attribution-level, and semantic-level interpretability for dense retrieval systems.
- We design a **bidirectional per-query modulation mechanism** enabling semantic alignment without re-encoding or expensive cross-attention.
- We develop a **token-level attribution method** using Moore-Penrose back-projection that causally links embedding modulations to human-interpretable keywords.
- We demonstrate **significant performance gains** across seven diverse benchmarks with minimal computational overhead.

2 Related Work

We organize related work into three categories: dense retrieval foundations, adapter mechanisms for performance enhancement, and the interpretability gap we address.

Dense Retrieval Methods. RAG systems employ a two-stage pipeline: an initial retriever selects candidates efficiently, followed by a reranker. Initial retrievers include lexical methods like BM25 (Robertson and Zaragoza, 2009b) with transparent term-matching, dense bi-encoders like DPR (Karpukhin et al., 2020) and Contriever (Izacard et al., 2022) capturing semantic similarity through learned embeddings, and hybrid approaches like SPLADE (Formal et al., 2021) combining contextual encoders with sparse representations. We focus on dense bi-encoders because they dominate modern RAG systems yet operate as black boxes.

Retrieval Adapters. Recent work has introduced lightweight adapters that enhance dense retrievers without retraining. *Architecture-modifying*

adapters (Rossi et al., 2024; Zeighami et al., 2025; Ding et al., 2023) update internal components, requiring white-box access. *Embedding-space adapters* operate on frozen encoder outputs: SEARCH-ADAPTOR (Yoon et al., 2024b), DIME (Campagnano et al., 2025b), and HYPENCODER (Killingback et al., 2025b) represent state-of-the-art embedding-space adaptation. *Embedding compression methods* (Liu et al., 2022b; Ma et al., 2021) reduce dimensionality but do not explain how transformations affect query-document interactions. *Parameter-efficient fine-tuning approaches* modify encoder parameters minimally: TART (Asai et al., 2023) and INSTRUCTOR (Su et al., 2023) inject task-specific instructions; LoRA (Hu et al., 2022) and IA3 (Liu et al., 2022a) update low-rank subspaces; PROMPTAGATOR (Dai et al., 2023) synthesizes training data via LLMs. We exclude these from our experimental comparison because they require encoder fine-tuning (violating the frozen-encoder constraint that enables IMRNNs plug-and-play deployment) and still provide no interpretability mechanisms. ADAPTED DENSE RETRIEVAL (Khatry et al., 2023) learns low-rank residuals for heterogeneous retrieval settings. Adapters have also been explored for sparse retrievers and rerankers (Hu et al., 2023), which fall outside our focus on dense initial retrieval. While these methods improve retrieval accuracy through various adaptation mechanisms, **none provide interpretability**: users cannot determine which semantic features drove retrieval decisions or how queries and documents are semantically aligned during ranking.

Interpretability in Dense retrievers. Interpretability for dense retrievers remain largely unaddressed, and existing approaches have critical limitations

- *Surrogate approximations* (Llordes et al., 2023b) fit sparse models to approximate dense rankings, but these post-hoc explanations may not faithfully represent the actual decision process, and approximation quality degrades as the sparse-dense gap widens.
- *Gradient-based attribution* (Yuksel and Kamps, 2025b) identifies high-gradient tokens during training but does not reveal semantic concepts emphasized during inference, and gradient explanations can be unstable (Adebayo et al., 2018).
- *Concept mapping* (Kang et al., 2025) aligns embedding dimensions with human-interpretable descriptors via sparse probing, but requires ad-

ditional annotation and does not explain query-document interactions during retrieval.

3 Problem Formulation

Task Definition: Given a user query q and a document corpus $\mathcal{D} = \{d_1, d_2, \dots, d_N\}$, the goal of the dense retriever is to rank all documents in \mathcal{D} by their relevance to q . A dense retrieval system consists of two components: (a) A **base encoder** $f_\theta : \mathcal{V} \rightarrow \mathbb{R}^n$ that maps text sequences (queries or documents) to fixed-dimensional embeddings, where \mathcal{V} is the space of all possible text sequences and n is the embedding dimension. (b) A **similarity function** $s : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ that computes relevance scores between query and document embeddings. Standard dense retrievers compute $\text{score}(q, d_i) = \cos(f_\theta(q), f_\theta(d_i))$, where $\cos(\cdot, \cdot)$ denotes cosine similarity. Documents are then ranked in descending order of these scores.

Current Limitation With Static Embeddings:

The embeddings $f_\theta(q)$ and $f_\theta(d_i)$ are *static*, computed independently and fixed after encoding. This means a document receives the same embedding regardless of which query it is being matched against, preventing the retriever from dynamically emphasizing query-relevant aspects of each document. Additionally, the system lacks interpretability: users cannot see which semantic dimensions drive the similarity score for a specific query-document pair.

3.1 Our approach: IMRNNs

IMRNNs address this limitation by introducing *dynamic modulation* on top of static embeddings from a frozen base encoder.

1. Dimension Reduction via Projection. To enable efficient learning and generalization, we project the high-dimensional base embeddings into a lower-dimensional working space. Let $\mathbf{P} \in \mathbb{R}^{m \times n}$ be a fixed linear projection matrix, where $m < n$ (we use $m = 256$ for $n = 1024$ in our experiments). For query q and document d_i , we compute:

$$\begin{aligned} \mathbf{q}_{\text{orig}} &= f_\theta(q) \in \mathbb{R}^n, & \mathbf{q}_{\text{proj}} &= \mathbf{P}\mathbf{q}_{\text{orig}} \in \mathbb{R}^m \\ \mathbf{d}_{\text{orig}}^{(i)} &= f_\theta(d_i) \in \mathbb{R}^n, & \mathbf{d}_{\text{proj}}^{(i)} &= \mathbf{P}\mathbf{d}_{\text{orig}}^{(i)} \in \mathbb{R}^m \end{aligned}$$

We optimize \mathbf{P} jointly with the adapters (described next), while the encoder f_θ remains frozen.

2. Query Adapter (\mathcal{A}_q) is a lightweight neural network that uses projected query embedding to produce a weight matrix and bias vector to modulate all document embeddings: $\mathcal{A}_q(\mathbf{q}_{\text{proj}}) = (\mathbf{W}_q \in \mathbb{R}^{m \times m}, \mathbf{b}_q \in \mathbb{R}^m)$. The modulated document embedding for document d_i is then computed as: $\mathbf{d}_{\text{mod}}^{(i)} = \mathbf{W}_q \cdot \mathbf{d}_{\text{proj}}^{(i)} + \mathbf{b}_q$. This transformation allows the query to *pull* semantically relevant documents closer in the embedding space and *push* irrelevant documents farther away, adapting document representations to the specific query context.

3. Document Adapter (\mathcal{A}_d) is a separate second lightweight neural network that processes each document embedding independently to produce document-specific transformations: $\mathcal{A}_d(\mathbf{d}_{\text{proj}}^{(i)}) = (\mathbf{W}_d^{(i)} \in \mathbb{R}^{m \times m}, \mathbf{b}_d^{(i)} \in \mathbb{R}^m)$. These transformations are aggregated across all documents to create a corpus-level adaptation signal:

$$\bar{\mathbf{W}}_d = \frac{1}{N} \sum_{i=1}^N \mathbf{W}_d^{(i)}, \quad \bar{\mathbf{b}}_d = \frac{1}{N} \sum_{i=1}^N \mathbf{b}_d^{(i)}$$

The modulated query embedding is then computed as: $\mathbf{q}_{\text{mod}} = \bar{\mathbf{W}}_d \cdot \mathbf{q}_{\text{proj}} + \bar{\mathbf{b}}_d$. This enables the query embedding to adapt to the characteristics of the document corpus, thereby aligning with the vocabulary and semantic space of the available documents.

4. Scoring and Ranking. After bidirectional modulation, the final relevance score between query q and document d_i is computed using cosine similarity:

$$\text{score}_{\text{IMRNNs}}(q, d_i) = \cos(\mathbf{q}_{\text{mod}}, \mathbf{d}_{\text{mod}}^{(i)})$$

Layer normalization is applied to both \mathbf{q}_{mod} and $\mathbf{d}_{\text{mod}}^{(i)}$ before computing cosine similarity to ensure stable gradients and bounded scores.

3.2 Training

IMRNNs are trained using a margin-based ranking loss over query-document pairs. For each query q in a training batch \mathcal{B} , we sample one relevant document d^+ and one irrelevant document d^- with BM25 for a hard negative example. The loss function is:

$$\mathcal{L} = \frac{1}{|\mathcal{B}|} \sum_{(q, d^+, d^-) \in \mathcal{B}} \max \left\{ 0, \gamma - \cos(\mathbf{q}_{\text{mod}}, \mathbf{d}_{\text{mod}}^+) + \cos(\mathbf{q}_{\text{mod}}, \mathbf{d}_{\text{mod}}^-) \right\}$$

where $\gamma > 0$ is a margin hyperparameter that enforces a minimum separation between relevant and irrelevant pairs in the modulated embedding space.

During training, only the parameters of \mathcal{A}_q and \mathcal{A}_d are updated while the base LLM encoder (such as e5-large, MiniLM, and BGE) (Wang et al., 2024), f_θ remain frozen. We optimize using the Adam optimizer with weight decay regularization to prevent overfitting. To bound computational cost and memory requirements, training operates on a subset of top- k BM25-retrieved candidates per query rather than the entire corpus. Early stopping is applied based on validation set performance to determine the optimal number of training epochs. Additional hyperparameter settings and implementation details are provided in Section 4.

3.3 Inference Workflow

At inference time, IMRNNs operate as follows:

1. For a new query q , compute $\mathbf{q}_{\text{orig}} = f_\theta(q)$ and $\mathbf{q}_{\text{proj}} = \mathbf{P}\mathbf{q}_{\text{orig}}$. For all documents in the corpus (pre-computed offline), obtain $\{\mathbf{d}_{\text{proj}}^{(i)}\}_{i=1}^N$.
2. Pass \mathbf{q}_{proj} through \mathcal{A}_q to obtain $(\mathbf{W}_q, \mathbf{b}_q)$, then compute modulated document embeddings:

$$\mathbf{d}_{\text{mod}}^{(i)} = \mathbf{W}_q \cdot \mathbf{d}_{\text{proj}}^{(i)} + \mathbf{b}_q \quad \text{for } i = 1, \dots, N$$

3. Pass each $\mathbf{d}_{\text{proj}}^{(i)}$ through \mathcal{A}_d to obtain $\{(\mathbf{W}_d^{(i)}, \mathbf{b}_d^{(i)})\}_{i=1}^N$, then aggregate and modulate the query:

$$\mathbf{q}_{\text{mod}} = \bar{\mathbf{W}}_d \cdot \mathbf{q}_{\text{proj}} + \bar{\mathbf{b}}_d$$

4. For each document d_i , compute $\text{score}_{\text{IMRNNs}}(q, d_i) = \cos(\mathbf{q}_{\text{mod}}, \mathbf{d}_{\text{mod}}^{(i)})$ and rank documents in descending order of score.

Computational Efficiency: The adapter networks \mathcal{A}_q and \mathcal{A}_d are lightweight (2-layer MLPs), adding minimal overhead. Document projections $\mathbf{d}_{\text{proj}}^{(i)}$ can be pre-computed offline and cached, so adapter forward passes dominate online cost per query and scale linearly with corpus size.

3.4 Interpretability Mechanism

The modulation framework enables analyzing the change induced in the original embedding space for direct interpretability. For any query-document pair (q, d_i) , we compute the modulation vectors:

$$\begin{aligned} \Delta \mathbf{q} &= \mathbf{q}_{\text{mod}} - \mathbf{q}_{\text{proj}} \\ \Delta \mathbf{d}^{(i)} &= \mathbf{d}_{\text{mod}}^{(i)} - \mathbf{d}_{\text{proj}}^{(i)} \end{aligned}$$

The change in retrieval score is measured as: $\Delta \text{similarity} = \cos(\mathbf{q}_{\text{mod}}, \mathbf{d}_{\text{mod}}^{(i)}) - \cos(\mathbf{q}_{\text{proj}}, \mathbf{d}_{\text{proj}}^{(i)})$

Positive Δ similarity indicates the modulation pulled the query-document pair closer (increasing relevance), while negative values indicate they

were pushed apart (decreasing relevance). To interpret these vectors in terms of semantic concepts, we back-project them to the original encoder’s embedding space using the Moore-Penrose pseudoinverse $\mathbf{P}^+ = \mathbf{P}^\top (\mathbf{P}\mathbf{P}^\top)^{-1}$ as follows:

$$\begin{aligned} \Delta \mathbf{q}_{\text{orig}} &= \mathbf{P}^+ \Delta \mathbf{q} \\ \Delta \mathbf{d}_{\text{orig}}^{(i)} &= \mathbf{P}^+ \Delta \mathbf{d}^{(i)}. \end{aligned}$$

In general, with $\mathbf{P} = \mathbf{U}\Sigma\mathbf{V}^\top$ (SVD), we can say that $\mathbf{P}^+ = \mathbf{V}\Sigma^+\mathbf{U}^\top$.

Let $\mathbf{E} \in \mathbb{R}^{|\mathcal{V}| \times n}$ denote the encoder’s token embedding table, where each row \mathbf{e}_t is the embedding of token $t \in \mathcal{V}$. We compute cosine similarity between the back-projected modulation vector and each token embedding:

$$\text{score}_t = \cos(\Delta \mathbf{q}_{\text{orig}}, \mathbf{e}_t) = \frac{\Delta \mathbf{q}_{\text{orig}}^\top \mathbf{e}_t}{\|\Delta \mathbf{q}_{\text{orig}}\|_2 \|\mathbf{e}_t\|_2}$$

Tokens with high positive scores indicate concepts that the modulation emphasized (pulling the embedding toward), while tokens with high negative scores indicate concepts that were de-emphasized (pushing away). By ranking tokens by $|\text{score}_t|$ and examining the top-ranked tokens, we obtain human-interpretable explanations of what semantic features drove the retrieval decision.

4 Experimental Setup

We evaluate IMRNNs on two complementary tasks across diverse datasets, comparing against state-of-the-art retrieval adapter baselines.

Retrieval Task. In the retrieval task, we determine the effectiveness of IMRNNs by comparing document retrieval accuracy with the base dense retriever and recent popular competing adaptation methods. Performance gain is evaluated on BEIR’s held-out test sets using MRR, nDCG, and Recall.

Interpretability Task. We qualitatively examine the semantic transparency of IMRNNs’ modulation mechanism. For selected queries, we analyze modulation vectors for both relevant and irrelevant documents, computing the Δ similarity (the change from the original to the modulated scores) and back-projecting the modulation vectors to identify top-ranked tokens using the Moore-Penrose pseudoinverse (Section 3). We validate interpretability by examining whether the identified keywords align with document relevance and whether changes in ranking correlate with meaningful semantic shifts.

Underlying Base Retrievers : We consider e5-large-v2, MiniLM-V6-L2 and BGE-Large-en as base retrievers and use it as baselines.

Dataset Selection Rationale. We selected seven datasets from the BEIR benchmark (Apache License 2.0) to ensure broad coverage in evaluating IMRNNs across multiple dimensions. MS MARCO tests scalability with large-scale web search (8.8M documents). Natural Questions and HotpotQA use Wikipedia, where HotpotQA requires multiple passages for multi-hop reasoning to synthesize answers, testing whether modulation connects semantically related but lexically distinct evidence. SciFact and TREC-COVID evaluate domain-specific terminology and precise semantic matching in scientific/biomedical retrieval. FiQA-2018 tests adaptation to financial jargon and numerical reasoning. Webis-Touché 2020 involves argumentation retrieval and depends on identifying viewpoints rather than topical overlap. Together, these span corpus sizes from 5K to 8.8M documents, single-hop to multi-hop reasoning, and general to highly specialized domains. More details are available in [Appendix A](#).

Baselines. We compare three adapter families:

- **DIME variants:** DIME variants represent query-only modulation via dimension selection. Select embedding subspaces by ranking dimensions via ℓ_2 magnitude. We evaluate 20%, 40%, 60%, and 80% dimensionality reduction to test whether simple dimension selection competes with learned modulation.
- **SEARCH-ADAPTOR:** Learns dataset-specific residual transformations applied globally to all queries and documents. This evaluates static, dataset-level adaptation.
- **HYPENCODER variants:** HYPENCODER variants represent query-only modulation via neural scoring. Generates query-specific MLPs to score documents, replacing cosine similarity with learned neural functions. We evaluate two, four, six, and eight hidden layers configurations to consider increasing expressive scoring functions.
- IMRNNs is our approach with query and document modulation.

All baselines operate on identical frozen encoder embeddings and use the same data splits. Implementation details and the computational efficiency comparison are provided in [Appendix B](#) and [Appendix C](#) respectively. Finally, we used Moore-

Penrose pseudoinverse to map modulation vectors back to actual vocabulary tokens.

5 Discussion on Interpretability

Interpretability Task. Table 1 shows that IMRNNs explain retrieval decisions through influential keywords that directly cause score changes. Consider the query `What currency is used in Mexico?` and the extracted keywords `{peso, Mexico}`. As shown in the Table 1 the similarity of the relevant document after modulation drops slightly from 0.94 to 0.87 ($\Delta = -0.07$) as the modulation strengthened the extracted keywords in the embedding thereby refining the match. In contrast, for the same query with keywords `{raining, sky}`, the score for an irrelevant document scores plummets significantly from 0.39 to 0.05 ($\Delta = -0.34$) as the modulation identified off-topic concepts and pushed the document away. Here the keywords explain a shift in the embeddings, bringing relevant concepts closer and pushing away from less relevant ones.

Similar pattern holds across other queries. For a query, `What county is incline village nv?`. Here the keywords `{Incline, Nevada}` identify the location terms that distinguish the relevant document from an irrelevant document with keywords `{Airport, Regional}`. As a result the score of the irrelevant documents drops from 0.41 to 0.11, as the recognized terms don't answer a location query. For the medical query about diabetic bleeding risk, the relevant document surfaces `{diabetes, coronary}` while the irrelevant document shows `{behavioural, metabolism}`. The IMRNNs identify which terms are relevant to the specific medical question.

The method identifies tokens whose embeddings align with the modulation direction; these are the semantic features that the adapters emphasized or suppressed. `peso` appears in the keyword list, clearly indicating that the modulation moved the embedding toward the dimension where `peso` lives in the vocabulary space. The ranking changes validate that keywords capture real semantic reasoning. In queries 1 and 3, the correct document doesn't appear in the original top-5 but appears in top-5 after modulation. Here the extracted keywords, `{Incline, Nevada}` and `{diabetes, coronary}`, are the key concepts. In query 2, the relevant document advances from rank 2 to rank 1, with `peso` as the decisive keyword. This indicates

Table 1: **Qualitative Examples of Modulation Effects on Retrieval Performance.** For each query, we show a relevant (green) and irrelevant (red) document pair. Modulation modestly reduces similarity for relevant documents ($\Delta \text{ Sim.} \approx -0.06$) while drastically reducing it for irrelevant ones ($\Delta \text{ Sim.} \approx -0.30$), improving discrimination. This enables relevant documents to rise in rankings (e.g., entering the top-5 as shown in the rightmost columns). Keywords extracted from queries and documents aid interpretability. **Green** indicates relevant; **red** indicates irrelevant.

Query	Document Type	Document Text	Orig. Sim.	Mod. Sim.	$\Delta \text{ Sim.}$	Query Key-words	Document Keywords	Orig Top-5 Doc. (Before)	Mod Top-5 Doc. (After)
What currency is used in Mexico?	Relevant	Mexico, Peso. The Mexican Peso is the currency of Mexico. Our currency rankings show that the most popular Mexico Peso exchange ...	0.94	0.87	-0.06	Mexican, Ridges	Airport, Peso, Mexico	1. No 2. Yes 3. No 4. No 5. No	1. Yes 2. No 3. No 4. No 5. No
	Irrelevant	<i>All you have to do is tune to the right channel or visit any number of weather and news Web sites and ...</i>	0.39	0.05	-0.34	Mexican, Ridges	Raining, Innate, Sky	-	-
What county is incline village nv?	Relevant	(Redirected from Incline Village-Crystal Bay, Nevada) Incline Village is a census-designated place in Washoe County, Nevada on the north shore ...	0.91	0.85	-0.06	Incline, Rupert	Incline, Hilly, Nevada	1. No 2. No 3. No 4. No 5. No	1. No 2. No 3. No 4. No 5. Yes
	Irrelevant	<i>Public Meetings: The Alva Regional Airport Commission will ...</i>	0.41	0.11	-0.29	Incline, Rupert	Airport, Regional, Sanctioned	-	-

that the system isn’t just re-scoring randomly as it identifies relevant semantic concepts for a given query and adjusts the embeddings accordingly.

Relevant documents show small similarity changes ($\Delta = -0.06$ to -0.08) because the modulation preserves already-good matches while refining them. Irrelevant documents show large drops ($\Delta = -0.29$ to -0.34) because the modulation actively suppresses mismatches. This asymmetry demonstrates principled behavior: it strengthens correct alignments and weakens incorrect ones.

6 Discussion on Retrieval Performance

Tables 2, 3, and 4 demonstrate that IMRNNs’ bidirectional modulation mechanism consistently improves ranking quality across diverse retrieval scenarios. The performance gains stem from how the adapters reshape the embedding space: the Query Adapter generates transformations that pull semantically relevant documents closer to the query while pushing irrelevant ones away, and the Document Adapter aggregates corpus-level signals to help the query align with the vocabulary and semantic struc-

ture of available documents. This dual adaptation addresses the core limitation of static embeddings, they cannot adjust to query-specific relevance signals or corpus-specific terminology. The nDCG rewards systems that place highly relevant documents at top positions with logarithmic discounting (Järvelin and Kekäläinen, 2002). By modulating embeddings to strengthen correct semantic alignments, IMRNNs ensures that the most relevant documents rise to ranks 1-3 where users actually look, rather than languishing at ranks 8-10 where they contribute little to user satisfaction.

SEARCH-ADAPTOR learns dataset-level residual transformations that apply uniformly across all queries, but struggles because different queries require different semantic adjustments, a geographic query about “incline village” needs location-term emphasis, while a medical query about “diabetic bleeding risk” needs disease-term emphasis. A single global transformation cannot capture this diversity, explaining why SEARCH-ADAPTOR often underperforms the base retriever. DIME variants progressively degrade as they remove more

Table 2: Results on Open-Domain Datasets (MS MARCO, Natural Questions, and HotpotQA).

Methods	MS MARCO			Natural Questions			HotpotQA			Average		
	nDCG	R	MRR	nDCG	R	MRR	nDCG	R	MRR	nDCG	R	MRR
e5-large-v2	0.85	0.97	0.81	0.72	0.90	0.68	0.54	0.61	0.65	0.70	0.83	0.71
MiniLM	0.91	0.97	0.90	0.89	0.68	0.73	0.61	0.64	0.53	0.80	0.76	0.72
BGE	0.88	0.97	0.89	0.73	0.90	0.69	0.61	0.64	0.54	0.74	0.84	0.71
DIME 20%	0.85	0.97	0.81	0.72	0.89	0.68	0.53	0.60	0.64	0.70	0.82	0.71
DIME 40%	0.84	0.97	0.80	0.72	0.90	0.67	0.51	0.58	0.62	0.69	0.82	0.70
DIME 60%	0.83	0.96	0.79	0.70	0.88	0.66	0.47	0.54	0.57	0.67	0.79	0.67
DIME 80%	0.80	0.93	0.75	0.64	0.83	0.60	0.37	0.42	0.45	0.60	0.73	0.60
SearchAd	0.84	0.96	0.80	0.71	0.89	0.67	0.37	0.43	0.46	0.64	0.76	0.64
Hyp (2)	0.68	0.74	0.71	0.43	0.60	0.40	0.29	0.25	0.27	0.47	0.53	0.46
Hyp (4)	0.69	0.74	0.70	0.44	0.60	0.40	0.29	0.26	0.27	0.47	0.53	0.46
Hyp (6)	0.71	0.75	0.71	0.44	0.61	0.40	0.30	0.26	0.28	0.48	0.54	0.46
Hyp (8)	0.70	0.75	0.72	0.45	0.61	0.42	0.30	0.26	0.29	0.48	0.54	0.48
IMRNNs (w e5-large-v2)	0.88	0.99	0.85	0.75	0.93	0.71	0.59	0.63	0.66	0.74	0.85	0.74
IMRNNs (w MiniLM)	0.93	0.99	0.92	0.91	0.69	0.75	0.63	0.66	0.56	0.82	0.78	0.74
IMRNNs (w BGE)	0.93	0.99	0.91	0.75	0.93	0.71	0.63	0.66	0.57	0.77	0.86	0.73

Table 3: Results on Domain-Specific Datasets (Scifact, Trec-COVID, and Webis-Touche2020).

Methods	Scifact			Trec-COVID			Webis-Touche			Average		
	nDCG	R	MRR	nDCG	R	MRR	nDCG	R	MRR	nDCG	R	MRR
e5-large-v2	0.72	0.87	0.68	0.79	0.03	1.00	0.60	0.31	0.85	0.70	0.40	0.84
MiniLM	0.79	0.63	0.66	0.75	0.02	0.87	0.59	0.30	0.85	0.62	0.31	0.65
BGE	0.67	0.84	0.66	0.76	0.02	0.92	0.56	0.29	0.85	0.62	0.38	0.71
DIME 20%	0.72	0.87	0.67	0.82	0.03	1.00	0.59	0.30	0.85	0.71	0.40	0.84
DIME 40%	0.70	0.87	0.65	0.82	0.03	1.00	0.61	0.32	0.85	0.71	0.41	0.83
DIME 60%	0.68	0.84	0.64	0.84	0.03	0.94	0.60	0.31	0.85	0.71	0.39	0.81
DIME 80%	0.60	0.78	0.55	0.77	0.03	0.86	0.62	0.32	0.78	0.66	0.38	0.73
SearchAd	0.69	0.85	0.65	0.84	0.03	1.00	0.58	0.31	0.85	0.70	0.40	0.83
Hyp (2)	0.62	0.75	0.63	0.78	0.02	0.59	0.59	0.29	0.83	0.66	0.35	0.68
Hyp (4)	0.61	0.63	0.68	0.76	0.01	0.44	0.59	0.28	0.83	0.65	0.31	0.65
Hyp (6)	0.72	0.62	0.60	0.79	0.01	0.47	0.60	0.29	0.84	0.70	0.31	0.64
Hyp (8)	0.65	0.81	0.68	0.84	0.01	0.52	0.60	0.29	0.85	0.70	0.37	0.68
IMRNNs (w e5-large-v2)	0.74	0.88	0.70	0.85	0.04	1.00	0.62	0.34	0.91	0.74	0.42	0.87
IMRNNs (w MiniLM)	0.82	0.65	0.69	0.78	0.03	1.00	0.61	0.31	0.90	0.65	0.35	0.68
IMRNNs (w BGE)	0.70	0.88	0.70	0.80	0.04	1.00	0.60	0.31	0.90	0.66	0.41	0.75

dimensions, demonstrating that magnitude-based selection discards information critical for semantic matching, a dimension that appears unimportant globally but may be essential for specific queries. HYPENCODER generates query-conditioned neural scorers, but these black-box functions lack the explicit semantic grounding that IMRNNs’ modulation vectors provide, and adding more layers (2→8) yields diminishing returns without addressing the fundamental need for interpretable semantic adaptation. The consistent pattern across open-domain datasets (MS MARCO, Natural Questions, HotpotQA) and specialized domains (SciFact, TREC-COVID, FiQA, Webis-Touché) confirms that IMRNNs’ approach generalizes: the mechanism adapts to whatever semantic features matter for each dataset, whether lexical overlap, multi-hop reasoning, domain terminology, or argumentative stance.

The FiQA anomaly (Table 4) illuminates a criti-

cal dependency: IMRNNs amplifies the quality of their base embeddings rather than replacing them. When e5-large-v2, MiniLM, and BGE embeddings are poorly calibrated for financial terminologies, IMRNNs built atop them cannot match HYPENCODER, which uses its own embedding generation. However, stacking IMRNNs on top of Hypencoder (8) embeddings yields substantial gains (9.09% nDCG, 4.76% Recall, 7.68% MRR), demonstrating that the modulation mechanism successfully enhances any sufficiently rich embedding space. This reveals IMRNNs’ architectural advantage: the adapters operate as a plug-and-play layer that improves whatever base retriever provides the best embeddings for a given domain, rather than requiring full model retraining or domain-specific architecture changes. The modulation vectors capture semantic refinements that static embeddings miss, emphasizing currency-related dimensions for financial queries, location dimensions for geo-

Table 4: Results on FiQA.

Methods	FiQA		
	nDCG	R	MRR
e5-large-v2	0.20	0.23	0.28
MiniLM	0.22	0.26	0.32
BGE	0.23	0.27	0.34
DIME 20%	0.20	0.22	0.28
DIME 40%	0.19	0.21	0.27
DIME 60%	0.18	0.20	0.25
DIME 80%	0.14	0.16	0.23
SearchAd	0.17	0.20	0.25
Hyp (2)	0.32	0.40	0.39
Hyp (4)	0.31	0.42	0.35
Hyp (6)	0.32	0.42	0.37
Hyp (8)	0.33	0.42	0.39
IMRNNs(e5-large-v2)	0.22	0.24	0.29
IMRNNs (w MiniLM)	0.23	0.27	0.33
IMRNNs (w BGE)	0.25	0.29	0.37
IMRNNs (Hyp(8))	0.36	0.44	0.42

graphic queries, and disease dimensions for medical queries, while preserving the computational efficiency of cosine similarity.

Comparing base retrievers, MiniLM, BGE and e5 with and without IMRNN, it is clear that with MiniLM-v6, IMRNNs yielded a 3.25% increase in MRR, 5.12% in Recall, and 3.64% in NDCG. These improvements were even more pronounced on BGE, where IMRNNs achieved gains of 5.2% in MRR, 4.4% in Recall, and 5.3% in NDCG. We observe that the magnitude of improvement scales with the model’s complexity as evident from comparing MiniLM, BGE and e5.

7 Conclusion

We introduce IMRNNs, the first lightweight retrieval adapters that make the embeddings of dense retrievers interpretable by achieving three levels of interpretability: structural, attribution, and semantic. We benchmark the retrieval performance of IMRNNs against state-of-the-art retrieval adapter baselines on diverse datasets and demonstrate that IMRNNs adapts query and document embeddings more effectively than competing methods, while also showing strong generalization. The semantic-level interpretability of IMRNNs is especially useful in applications where access to key tokens or keywords in both the query and the documents plays a major role. We release all code and scripts under the CC BY 4.0 license for reproducibility.¹

¹<https://github.com/YashSaxena21/IMRNNs>

8 Limitations

Experiments reveal three key limitations of IMRNNs. First, the token-level attribution method can produce noisy mappings where some identified tokens (e.g., “Ridges”, “Innate” in Table 1) appear semantically unclear or spuriously correlated with the query-document relationship. This occurs because back-projecting continuous embedding modulations to discrete tokens via pseudoinverse is inherently approximate, and the closest token in embedding space may not semantically correspond to the actual concept driving the modulation. Filtering is often necessary to obtain interpretable explanations, but systematic methods for identifying spurious tokens remain an open challenge. Second, IMRNNs incur higher inference latency than dimension-selection methods (Table 6 in Appendix) because the Document Adapter must process each corpus document individually to compute transformations. While this bidirectional modulation enables richer semantic adaptation, it scales linearly with corpus size, potentially limiting deployment on extremely large corpora (>10M documents) without infrastructure optimizations like caching or approximate nearest neighbor filtering. Third, IMRNNs amplify rather than replace the quality of base embeddings. When base retrievers produce poorly calibrated embeddings for a domain (e.g., financial terminology in FiQA), IMRNNs cannot compensate for fundamental semantic gaps. The modulation mechanism assumes the base embedding space already captures relevant semantic dimensions, it refines their emphasis rather than introducing new concepts. This dependency suggests IMRNNs are best deployed as an enhancement layer atop domain-appropriate base retrievers rather than a universal solution.

Acknowledgements

We thank the ACL ARR reviewers for their constructive feedback that significantly improved this work. We are grateful to Mandar Chaudhary and the students in the Knowledge-infused AI and Inference Lab at UMBC for their insightful discussions and reviews. This work was supported in part by USISTEF and the UMBC Cybersecurity Initiative. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies of the funding agencies.

References

- Julius Adebayo, Justin Gilmer, Michael Muehly, Ian Goodfellow, Moritz Hardt, and Been Kim. 2018. Sanity checks for saliency maps. *Advances in Neural Information Processing Systems*, 31.
- Dustin Arendt, Zhuanyi Shaw, Prasha Shrestha, Ellyn Ayton, Maria Glenski, and Svitlana Volkova. 2021. Crosscheck: Rapid, reproducible, and interpretable model evaluation. In *Proceedings of the Second Workshop on Data Science with Human in the Loop: Language Advances*, pages 79–85.
- Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi, and Wen-tau Yih. 2023. [Task-aware retrieval with instructions](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3650–3675, Toronto, Canada. Association for Computational Linguistics.
- João Carlos Alves Barata and Mahir Saleh Hussein. 2012. The moore–penrose pseudoinverse: A tutorial review of the theory. *Brazilian Journal of Physics*, 42(1):146–165.
- Nitay Calderon and Roi Reichart. 2025. On behalf of the stakeholders: Trends in nlp model interpretability in the era of llms. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 656–693.
- Cesare Campagnano, Antonio Mallia, and Fabrizio Silvestri. 2025a. Unveiling dime: Reproducibility, generalizability, and formal analysis of dimension importance estimation for dense retrieval. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3367–3376.
- Cesare Campagnano, Antonio Mallia, and Fabrizio Silvestri. 2025b. [Unveiling dime: Reproducibility, generalizability, and formal analysis of dimension importance estimation for dense retrieval](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '25*, page 3367–3376, New York, NY, USA. Association for Computing Machinery.
- Zhuyun Dai, Vincent Y Zhao, Ji Ma, Yi Luan, Jianmo Ni, Jing Lu, Anton Bakalov, Kelvin Guu, Keith Hall, and Ming-Wei Chang. 2023. [Promptagator: Few-shot dense retrieval from 8 examples](#). In *The Eleventh International Conference on Learning Representations*.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, and 1 others. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature machine intelligence*, 5(3):220–235.
- Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. [SPLADE: sparse lexical and expansion model for first stage ranking](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 2288–2292. ACM.
- Michael Glass, Gaetano Rossiello, Md Faisal Mahub Chowdhury, Ankita Naik, Pengshan Cai, and Alfio Gliozzo. 2022. Re2g: Retrieve, rerank, generate. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2701–2715.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Zhiqiang Hu, Lei Wang, Yihui Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya Poria, and Roy Lee. 2023. [LLM-adapters: An adapter family for parameter-efficient fine-tuning of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5254–5276, Singapore. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Transactions on Machine Learning Research*.
- Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of ir techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4):422–446.
- Hao Kang, Tevin Wang, and Chenyan Xiong. 2025. [Interpret and control dense retrieval with sparse latent features](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 700–709, Albuquerque, New Mexico. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Anirudh Khatry, Yasharth Bajpai, Priyanshu Gupta, Sumit Gulwani, and Ashish Tiwari. 2023. [Augmented embeddings for custom retrievals](#). *Preprint*, arXiv:2310.05380.
- Julian Killingback, Hansi Zeng, and Hamed Zamani. 2025a. Hypencoder: Hypernetworks for information

- retrieval. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2372–2383.
- Julian Killingback, Hansi Zeng, and Hamed Zamani. 2025b. [Hypencoder: Hypernetworks for information retrieval](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '25, page 2372–2383, New York, NY, USA. Association for Computing Machinery.
- Sejong Kim, Hyunseo Song, Hyunwoo Seo, and Hyunjun Kim. 2025. Optimizing retrieval strategies for financial question answering documents in retrieval-augmented generation systems. *arXiv preprint arXiv:2503.15191*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474.
- Yizhi Li, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2021. More robust dense retrieval with contrastive dual learning. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 287–296.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohata, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022a. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Zhenghao Liu, Han Zhang, Chenyan Xiong, Zhiyuan Liu, Yu Gu, and Xiaohua Li. 2022b. [Dimension reduction for efficient dense retrieval via conditional autoencoder](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5692–5698, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Michael Llorde, Debasis Ganguly, Sumit Bhatia, and Chirag Agarwal. 2023a. [Explain like i am bm25: Interpreting a dense model's ranked-list with a sparse approximation](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1976–1980.
- Michael Llorde, Debasis Ganguly, Sumit Bhatia, and Chirag Agarwal. 2023b. [Explain like i am bm25: Interpreting a dense model's ranked-list with a sparse approximation](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '23, page 1976–1980, New York, NY, USA. Association for Computing Machinery.
- Xueguang Ma, Minghan Li, Kai Sun, Ji Xin, and Jimmy Lin. 2021. [Simple and effective unsupervised redundancy elimination to compress dense vectors for passage retrieval](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2854–2859, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D Manning, and Daniel E Ho. 2025. Hallucination-free? assessing the reliability of leading ai legal research tools. *Journal of Empirical Legal Studies*, 22(2):216–242.
- Monica Munnangi, Akshay Swaminathan, Jason Alan Fries, Jenelle Jindal, Sanjana Narayanan, Ivan Lopez, Lucia Tu, Philip Chung, Jesutofunmi A Omiye, Mehr Kashyap, and 1 others. 2025. Factehr: A dataset for evaluating factuality in clinical notes using llms.
- Dheeraj Rajagopal, Vidhisha Balachandran, Eduard H Hovy, and Yulia Tsvetkov. 2021. Selfexplain: A self-explaining architecture for neural text classifiers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 836–850.
- Stephen Robertson and Hugo Zaragoza. 2009a. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Stephen Robertson and Hugo Zaragoza. 2009b. [The probabilistic relevance framework: Bm25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Nicholas Rossi, Juexin Lin, Feng Liu, Zhen Yang, Tony Lee, Alessandro Magnani, and Ciya Liao. 2024. [Relevance filtering for embedding-based retrieval](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, page 4828–4835, New York, NY, USA. Association for Computing Machinery.
- Hassan Sajjad, Nadir Durrani, and Fahim Dalvi. 2022. Neuron-level interpretation of deep nlp models: A survey. *Transactions of the Association for Computational Linguistics*, 10:1285–1303.
- Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2023. [One embedder, any task: Instruction-finetuned text embeddings](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1102–1121, Toronto, Canada. Association for Computational Linguistics.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2023. [Interpretability in the wild: a circuit for indirect object identification in GPT-2 small](#). In *The Eleventh International Conference on Learning Representations*.

- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2024. [Text embeddings by weakly-supervised contrastive pre-training](#). *Preprint*, arXiv:2212.03533.
- Jinsung Yoon, Yanfei Chen, Sercan Arik, and Tomas Pfister. 2024a. Search-adaptor: Embedding customization for information retrieval. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12230–12247.
- Jinsung Yoon, Yanfei Chen, Sercan Arik, and Tomas Pfister. 2024b. [Search-adaptor: Embedding customization for information retrieval](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12230–12247, Bangkok, Thailand. Association for Computational Linguistics.
- Goksenin Yuksel and Jaap Kamps. 2025a. Interpretability analysis of domain adapted dense retrievers. *arXiv preprint arXiv:2501.14459*.
- Goksenin Yuksel and Jaap Kamps. 2025b. [Interpretability analysis of domain adapted dense retrievers](#). *Preprint*, arXiv:2501.14459.
- Sepanta Zeighami, Zac Wellmer, and Aditya Parameswaran. 2025. [NUDGE: Lightweight non-parametric fine-tuning of embeddings for retrieval](#). In *The Thirteenth International Conference on Learning Representations*.
- Sheng Zhang, Jin Wang, Haitao Jiang, and Rui Song. 2022. Locally aggregated feature attribution on natural language model understanding. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2189–2201.
- Yilun Zhou, Serena Booth, Marco Tulio Ribeiro, and Julie Shah. 2022. Do feature attribution methods correctly attribute features? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9623–9633.
- Yujia Zhou, Yan Liu, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Zheng Liu, Chaozhuo Li, Zhicheng Dou, Tsung-Yi Ho, and Philip S Yu. 2024. Trustworthiness in retrieval-augmented generation systems: A survey. *arXiv preprint arXiv:2409.10102*.

A Dataset Selection

BEIR contains 15 retrieval datasets that cover a wide range of domains and query types. In the main paper, we selected seven datasets that provide sufficient diversity for evaluating both retrieval effectiveness and the interpretability mechanisms introduced by IMRNNs. The goal was to include datasets that differ meaningfully in domain, reasoning structure, and retrieval difficulty, rather than to exhaustively evaluate on all BEIR tasks.

Our chosen subset covers three important axes of variation:

- **Domain diversity.** The selected datasets span open-domain retrieval (MS MARCO, NQ), multi-hop reasoning (HotpotQA), scientific fact checking (Scifact), legal and argument-focused retrieval (Webis-Touche), biomedical retrieval (Trec-COVID), and financial question answering (FiQA). This includes all major retrieval settings that stress different aspects of semantic matching.
- **Query complexity.** We include datasets that require single-hop retrieval, multi-hop synthesis, and fact verification. This variation is important because the modulation mechanism in IMRNNs adapts embeddings based on both query semantics and corpus structure, which is not tested by simple single-hop retrieval alone.
- **Dataset size and structural properties.** The selected datasets cover a wide range of corpus sizes and document distributions. This ensures that IMRNNs is evaluated under different levels of retrieval sparsity, redundancy, and noise. These conditions affect how much semantic refinement the adapters can provide.

Several of the remaining BEIR datasets are stylistic variants of tasks already included or add little new retrieval structure relative to the selected subset. This selection approach is consistent with many recent retrieval studies that report results on a representative subset of BEIR when the full benchmark is not required to evaluate the proposed contribution. Our method is dataset-agnostic because interpretability in IMRNNs arises from the mathematical properties of modulation and back-projection, rather than dataset-specific lexical patterns. Adding more datasets would increase vol-

Table 5: BEIR Dataset Details.

Dataset	Domain	Type	#Queries	#Documents
MS MARCO	Web Search	Single-hop	6,980	8,841,823
Natural Questions	Wikipedia	Single-hop	3,452	2,681,468
HotpotQA	Wikipedia	Multi-hop	7,405	5,233,329
SciFact	Scientific	Single-hop	300	5,183
TREC-COVID	Biomedical	Single-hop	50	171,332
Webis-Touché 2020	Argumentation	Single-hop	49	382,545
FiQA-2018	Finance	Single-hop	648	57,638

ume but would not change the qualitative insights or the interpretability analysis.

B Implementation Details

IMRNNs use projection dimension $m = 256$. Both adapters are two-layer MLPs with ReLU activations and layer normalization. Training uses Adam (learning rate 10^{-4} , weight decay 10^{-5} , batch size 32) with margin $\gamma = 0.3$, operating on top-100 BM25 candidates per query. Early stopping uses patience of 5 epochs on validation nDCG. Training converges within 10-20 epochs on a single NVIDIA H100 GPU.

C Computational Time Comparison

Latency is averaged per single query, and throughput is computed as the inverse of average latency.

Table 6: Inference efficiency comparison (averaged over 1,000 queries).

Method	Latency (ms/query)	Throughput (queries/s)
Magnitude DIME	0.96	61.22
Search-Adaptor	1.02	58.82
Hypencoder	1.71	35.09
IMRNNs	1.64	36.59

D Additional Experiments

The additional experiments provided in Table 7 strengthen the main claim that IMRNNs is retriever-agnostic at the same time dataset independent and can be attached to a wide range of dense encoders without retraining them. Two observations are consistent across all models.

- The magnitude of improvement grows with the capacity of the base retriever. MiniLM performs the worst among the three base models, and its gains are smaller, while BGE performs better and shows larger improvements.
- Even compact models benefit from modulation. The adapters consistently improve the

Table 7: **Additional Experiments on ArguAna, Quora, and Scidocs datasets from the BEIR benchmark suite.**

Methods	ArguAna			Quora			Scidocs			Average		
	nDCG	R	MRR									
e5-large	0.78	0.94	0.87	0.90	0.89	0.96	0.25	0.29	0.95	0.64	0.70	0.92
MiniLM	0.86	0.72	0.75	0.88	0.84	0.92	0.15	0.22	0.89	0.63	0.59	0.85
BGE	0.75	0.92	0.88	0.89	0.84	0.90	0.20	0.27	0.92	0.61	0.67	0.90
IMRNNs (e5-large)	0.81	0.96	0.92	0.91	0.92	0.98	0.29	0.34	0.96	0.67	0.74	0.95
IMRNNs (MiniLM)	0.90	0.76	0.79	0.91	0.92	0.99	0.19	0.25	0.91	0.66	0.64	0.89
IMRNNs (BGE)	0.79	0.96	0.90	0.92	0.89	0.94	0.27	0.29	0.95	0.66	0.71	0.93

separation between relevant and irrelevant documents regardless of the dimensionality of the embedding space.

Overall, the results provide strong evidence that IMRNNs generalizes well beyond the specific encoder used in the main paper, and can be reliably deployed across diverse retrieval systems.