

Scaling Laws for Embedding Dimension in Information Retrieval

Julian Killingback*

jkillingback@cs.umass.edu

Center for Intelligent Information Retrieval
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Madine Manas

mmadine@cs.umass.edu

Center for Intelligent Information Retrieval
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Mahta Rafiee*

mrafiee@cs.umass.edu

Center for Intelligent Information Retrieval
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Hamed Zamani

zamani@cs.umass.edu

Center for Intelligent Information Retrieval
University of Massachusetts Amherst
Amherst, Massachusetts, USA

Abstract

Dense retrieval, which encodes queries and documents into a single dense vector, has become the dominant neural retrieval approach due to its simplicity and compatibility with fast approximate nearest neighbor algorithms. As the tasks dense retrieval performs grow in complexity, the fundamental limitations of the underlying data structure and similarity metric—namely vectors and inner-products—become more apparent. Recent work has shown theoretical limitations inherent to single vectors and inner-products, which are generally tied to the embedding dimension. Given the importance of embedding dimension for retrieval capacity, understanding how dense retrieval performance changes as embedding dimension is scaled is fundamental to building next-generation retrieval models that balance effectiveness and efficiency. In this work, we conduct a comprehensive analysis of the relationship between embedding dimension and retrieval performance. Our experiments include a range of model sizes from two model families to construct a detailed picture of embedding scaling behavior. We find that the scaling behavior fits a power law, allowing us to derive scaling laws for performance given only the embedding dimension, as well as a joint law accounting for embedding dimension and model size. Our analysis shows that for evaluation tasks aligned with the training task, performance continues to improve as embedding size increases, though with diminishing returns. For evaluation data that are less aligned with the training task, we find that performance is less predictable, with performance degrading with larger embedding dimensions for certain tasks. We hope our work provides additional insight into the limitations of embeddings and their behavior as well as offers a practical guide for selecting model and embedding dimension to achieve optimal performance with reduced storage and compute costs.

CCS Concepts

• **Information systems** → **Retrieval models and ranking; Data encoding and canonicalization; Search engine architectures and scalability.**

Keywords

Scaling Laws, Dense Retrieval, Neural Ranking, Embedding Dimension

ACM Reference Format:

Julian Killingback*, Mahta Rafiee*, Madine Manas, and Hamed Zamani. 2026. Scaling Laws for Embedding Dimension in Information Retrieval. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, July 20–24, 2026, Melbourne, VIC, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3805712.3809708>

1 Introduction

Dense retrieval has emerged as a cornerstone of modern neural information retrieval (IR) due to its effectiveness and simplicity. By encoding queries and documents into a shared latent space, bi-encoder architectures [15, 36] enable efficient retrieval over massive corpora using maximum inner product search (MIPS). However, the scope of dense retrieval has recently expanded beyond simple topical matching to include more sophisticated tasks such as instruction following and multi-step reasoning. As these tasks grow in complexity, the inherent limitations of the underlying data structure, i.e., fixed-dimensional vectors, become increasingly important.

Previous work has shown that the capacity of a dense retriever is fundamentally constrained by the geometry of its embedding space. Theoretical analyses have demonstrated that the ability of a model to perfectly separate relevant documents from a non-relevant set is strictly bounded by the embedding dimension d [17, 31]. For a corpus of size $M > d + 1$, there exist relevance patterns that cannot be represented by any linear similarity function. Despite these theoretical insights, the empirical relationship between embedding dimensionality and retrieval effectiveness has not been systematically quantified in the same way that model parameters or training data have been characterized through scaling laws [6, 14].

Though exploring the scaling behavior of embedding dimensions provides useful insight into existing theories, understanding how

* Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR '26, Melbourne, VIC, Australia*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2599-9/2026/07

<https://doi.org/10.1145/3805712.3809708>




Figure 1: Scaling behavior of contrastive entropy relative to embedding dimension on TREC DL Combined and MSMARCO Dev. The points represent the observed contrastive entropy while the line represents the fitted dimension-only scaling law. The BERT model shown in both plots is BERT-L8-H512-A8. The ETTIN model shown in both plots is Ettin L19-H512-A8.

retrieval performance scales with embedding dimension is not just a theoretical exercise; it is a practical tool for system design, as demonstrated by the impact of prior scaling laws [8, 14]. As an example of the impact scaling laws can have, the Chinchilla scaling laws proposed by Hoffmann et al. [8] showed that for a limited compute budget, an optimally trained model has both compute and data scaled equally. This finding allowed the Chinchilla model to outperform the existing Gopher model with $\approx 20\%$ of the parameters and generally empowered more efficient large language model training. Additionally, scaling laws enable experiments to be done at a small scale while being able to directly predict the outcome on a scaled-up version.

As shown by the impact of prior scaling laws, understanding and predicting how different parameters impact performance are essential tools for building effective and efficient models. For dense retrieval efficiency, the choice of embedding dimension d directly dictates the storage footprint of the document index and the computational cost of retrieval. These concerns are especially acute for on-device settings where disk space is often at a premium. Despite the impact of embedding dimension on efficiency, practitioners often rely on the “native” hidden size of a transformer backbone (e.g., 768 for BERT-base [4]), though it is unclear whether this choice is optimal or how performance is impacted as dimensions are expanded or compressed. Furthermore, the interplay between model size (parameters) and embedding size (representation capacity) remains poorly understood, leaving architects without a clear guide for balancing the two under fixed compute or storage budgets.

In this work, we conduct a comprehensive empirical study of embedding dimension scaling across two distinct model families: the classic BERT architecture [4] and the more recent ETTIN suite [32]. By training models across a wide range of parameter counts and embedding dimensions, including dimensions that exceed the backbone’s hidden size, we provide a detailed picture of retrieval scaling behavior. Our analysis reveals that retrieval performance, measured by contrastive entropy, follows a predictable power law relative to the embedding dimension. We further derive a joint scaling law that accounts for both model size and dimensionality, offering a unified framework for predicting performance.

Our findings show that for retrieval tasks aligned with the training objective, increasing the embedding dimension yields consistent, albeit diminishing, returns. However, we observe that this scaling behavior is less consistent for out-of-domain tasks which




Figure 2: Empirical results and joint scaling laws for the BERT model family on MSMARCO Dev and TREC DL Combined. The points represent empirical results at various embedding dimensions and model sizes which are represented by the point colors. The dashed lines represent the joint scaling laws fit on the observed data.

differ from the training distribution, though generally increasing the embedding dimension still improves performance. Finally, we demonstrate the utility of our scaling laws by providing a cost-aware analysis, showing the optimal model and embedding size to maximize retrieval effectiveness under specific inference compute constraints. We hope our work serves as a practical guide for the development of next-generation retrieval systems that are both effective and efficient.

2 Related Work

Theoretical Impact of Embedding Dimension. Theoretical analysis highlights the geometric bottlenecks of single-vector dense retrieval, where the capacity to model relevance is fundamentally tied to dimensionality [20, 21]. Killingback et al. [17] apply Radon’s




Figure 3: Empirical results and joint scaling laws for the ETTin model family on MSMARCO Dev and TREC DL Combined. The points represent empirical results at various embedding dimensions and model sizes which are represented by the point colors. The dashed lines represent the joint scaling laws fit on the observed data.

theorem to demonstrate that fixed-dimensional spaces cannot linearly separate arbitrary relevant sets once the corpus size exceeds the dimension plus one. Similarly, Weller et al. [31] prove that the number of realizable top- k result sets is strictly bounded by dimensionality, a limitation which they show can happen even with simple natural language queries. These findings suggest an inherent expressivity ceiling for inner-product dense retrieval that is tied to the embedding dimension, which motivates our systematic evaluation of how embedding dimension dictates performance.

Dense Retrieval and Efficient Representations. Dense retrieval models, such as DPR [15], encode text into fixed vectors, with effectiveness significantly improved via hard negative selection [25, 26, 34], distillation [9], and self-supervised pretraining [7, 11]. However, the embedding dimension—a primary driver of storage and latency for both exhaustive and approximate nearest neighbor (ANN) systems like Faiss [5]—is typically treated as a static parameter derived from the encoder’s hidden size. There has been some work on minimizing the size of the embedding dimension via compression or dimensionality reduction [20, 35], but these approaches generally operate within or below the bounds of the backbone’s hidden dimension. Matryoshka Representations [18] are similar, but allow for embedding truncation after encoding. Other approaches like product quantization [12] can help reduce embedding space as well. Our work is complementary to many of the explicit compression techniques and efficient search techniques which reduce the compute cost of retrieval, as we consider possibly reducing the dimension of the embeddings for added efficiency. In

contrast to existing approaches to reduce the embedding dimension, we treat the target embedding dimension as a controlled variable and systematically study both compression and expansion relative to the encoder hidden size, including regimes larger than the backbone. Our comprehensive analysis allows us to derive scaling laws that accurately predict the retrieval performance given different embedding dimensions which has not been possible before this point.

Scaling Laws for Retrieval. Scaling laws have successfully characterized performance relative to parameters, data, and compute in language modeling [8, 14] and dense retrieval [6, 24]. Zeng et al. expanded dense-only scaling behavior to compared scaling trends across sparse versus dense paradigms. Though scaling laws exist for dense retrieval, all prior work treats the embedding dimension as a fixed hyperparameter. In this work we specifically isolate the embedding dimension to understand how it contributes to retrieval performance as well as derive a joint scaling law for embedding dimension and model size.

3 Methodology

In this section, we describe the experimental setup we use to understand the scaling behavior of embedding dimensions. Our setup is largely similar to prior work in terms of model architecture and training approach to make our findings as realistic and directly applicable as possible. The few times we diverge from common training and architectural approaches are in service of our core research question, which is to understand how embedding scaling impacts retrieval performance.

3.0.1 Base Encoders. To get a complete picture of the scaling behavior of embedding dimensions, we use two encoder model families, BERT [4] and ETTin [32]. BERT represents one of the most popular encoder models while ETTin is a more recent suite and thus adopts many advancements in model architecture and training recipe. Both model families include a range of models trained on the same data making comparisons between models of different sizes possible. As model size plays a crucial role in both computational cost and retrieval performance, we select several model sizes across both families to understand the interplay between scaling model size and scaling embedding size. The specific models used along with their parameter counts are presented in Table 1.

The two distinct architectures ensure that our results generalize beyond a specific model configuration. To further bolster our results, each architecture uses a different training recipe based on popular dense retrieval approaches. We believe that this is an important choice for understanding the generalization of our findings and enabling a more complete analysis of how different training choices impact scaling behavior.

3.1 Model Architecture

Following existing dense retrieval approaches, we use a bi-encoder approach which uses a shared transformer encoder E to encode queries and documents [28]. Formally, given text T with tokens t_1, t_2, \dots, t_n , we use E to encode the text into hidden states h_1, h_2, \dots, h_n where $h_i \in \mathbb{R}^d$ and d is the hidden size of E . We then scale these

Table 1: Information about the models used in our experiments and the parameters of fitting our dimension only scaling law on MSMARCO Dev and TREC DL Combined.

Model Name	Size	Embed Dims	MSMARCO Dev				TREC DL			
			A	α	δ_D	R^2	A	α	δ_D	R^2
BERT-L2-H128-A2	4.39 M	32-2048	84.73	1.71	0.128	0.9988	319.19	1.73	0.953	0.9984
BERT-L12-H128-A2	6.37 M	32-2048	61.16	1.75	0.092	0.9991	72.49	1.34	0.810	0.9905
BERT-L4-H256-A4	11.17 M	32-4096	117.48	1.86	0.052	0.9962	63.14	1.21	0.517	0.9993
BERT-L4-H512-A8	28.76 M	32-8192	62.52	1.72	0.027	0.9995	105.71	1.35	0.394	0.9992
BERT-L2-H768-A12	38.60 M	48-12288	472.56	2.23	0.026	0.9996	179.78	1.48	0.359	0.9989
BERT-L8-H512-A8	41.37 M	32-8192	55.21	1.76	0.024	0.9998	158.77	1.48	0.341	0.9997
BERT-L12-H768-A12	109.48 M	48-12288	170.38	2.08	0.018	0.9917	96.42	1.33	0.342	0.9953
Ettin-L7-H256-A4	16.80 M	64-4096	3.46	1.05	0.149	0.9771	9.04	0.91	0.977	0.9782
Ettin-L10-H384-A6	31.88 M	96-6144	58.68	1.64	0.097	0.9806	4.47	0.75	0.769	0.9802
Ettin-L19-H512-A8	68.14 M	128-8192	1.75	0.91	0.058	0.9953	6.00	0.75	0.690	0.9941
Ettin-L22-H768-A12	149.01 M	192-12288	0.83	0.73	0.042	0.9857	11.27	0.80	0.598	0.9923
Ettin-L28-H1024-A16	394.78 M	256-16384	0.38	0.58	0.024	0.9467	14.79	0.87	0.558	0.9709
Ettin-L28-H1792-A28	1028.05 M	448-28672	0.06	0.45	0.022	0.5325	72.94	1.05	0.479	0.9187

values using a linear layer:

$$h'_i = Wh_i + b$$

where $W \in \mathbb{R}^{m \times d}$ represents the weight matrix and $b \in \mathbb{R}^m$ represents the bias vector where m represents the final embedding dimension. To get the final embedding e , we take the element-wise mean over all the hidden dimensions in the sequence.

$$e = \frac{1}{n} \sum_{i=1}^n h'_i$$

We selected mean pooling over other pooling methods as it has a smaller bottleneck than selection-based pooling methods such as selecting the start token or end token. For Ettin models, we normalize the final embeddings and apply a fixed temperature of $\tau = 0.02$ as larger hidden embedding dimensions used by Ettin can cause unnormalized dot-product similarities to grow with dimension, leading to inflated similarity values and numerical instability.

3.2 Training Data

To fully utilize the underlying architecture of the model families and generalize our results beyond specific training data, we use two different training datasets for the two different encoder families. For the BERT-based models, we use the MSMARCO Passage dataset [22] for training. Specifically, we use the version of MSMARCO preprocessed with teacher scores from Killingback et al. [17]. The dataset contains approximately 500k training queries. Each training instance consists of: (1) a query, (2) one or more known relevant passages (from MSMARCO annotations), and (3) around 200 additional candidate passages. All query-passage pairs include a score from the cross-encoder model cross-encoder/msmarco-MiniLM-L-12-v2 to provide teacher scores for training.

For the Ettin models, we use MSMARCO Instruct [22]. This dataset augments the standard MSMARCO Passage dataset with LLM-generated instructions to the original queries. Additional hard negatives are also generated to encourage instruction following. As the instruction-augmented queries are in addition to the existing MSMARCO queries, this dataset has approximately 1M training

queries. We selected this dataset as prior work suggested that it produced robust performance on complex retrieval tasks [16] and we wanted to understand whether longer and more complex queries would have an impact on scaling behavior.

During training, we sample one positive and seven negative passages per query. For the BERT-based model we selected the most relevant passage, as judged by the teacher, as the positive. For the MSMARCO Instruct data we prioritize sampling the up to three negatives which were generated to ensure the models are able to understand instructions, and the rest are sampled from query negatives. It is worth noting that for the BERT data, although we call the passages negative, following prior work, the passages might be relevant as they are sampled from retrieved documents for the query.

3.3 Training Objective

Our primary training recipe for the BERT models combines two loss functions, a knowledge distillation and contrastive learning loss. We made this choice to reflect current popular training practices which often combine the two for improved performance [17, 19, 27, 37]. For MSMARCO Instruct, we only use contrastive learning loss as the dataset does not have labels, and standard cross-encoders may not be well suited for the instruction following task. Additionally, this is also a common training setup for dense retrieval, so it is worth representing in our experiments.

3.3.1 Margin Mean Squared Error. The Margin MSE loss [10] encourages the encoder to reproduce the differences in score margin from a teacher cross-encoder. For a query q , positive passage embedding p^+ , and negative passage embedding p^- , the loss is:

$$\mathcal{L}_{\text{margin-mse}} = (\Delta_{\text{student}} - \Delta_{\text{teacher}})^2 \quad (1)$$

where $\Delta_{\text{student}} = (e_q \cdot e_{p^+}) - (e_q \cdot e_{p^-})$ and $\Delta_{\text{teacher}} = s_{\text{teacher}}(q, p^+) - s_{\text{teacher}}(q, p^-)$. The values e_q , e_{p^+} , and e_{p^-} represent the encoded query, positive passage, and negative passage, respectively, and s_{teacher} denotes the teacher cross-encoder scores. We apply Margin-MSE to each query and hard negative pair in the batch and take the mean to get the final loss.

3.3.2 Cross-Entropy Contrastive Loss. The contrastive loss treats retrieval as a classification problem over passages in each batch. For a query q with positive passage p^+ and negative passages \mathcal{N} that vary for the BERT training and ETTIN training, the contrastive loss is defined as:

$$\mathcal{L}_{\text{contrastive}} = -\log \frac{\exp(e_q \cdot e_{p^+})}{\exp(e_q \cdot e_{p^+}) + \sum_{p^- \in \mathcal{N}} \exp(e_q \cdot e_{p^-})} \quad (2)$$

For BERT, we do not include the seven hard negatives in \mathcal{N} as the sampling strategy results in relevant passages in the hard negatives; instead, we use only positive passages for the other in-batch queries. For ETTIN, we use all in-batch passages except p^+ . Additionally, for the ETTIN models we use a fixed temperature of $\tau = 0.02$ to account for the limited range due to the embedding normalization.

3.3.3 Combined Objective. The final training loss for BERT is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{margin-mse}} + \mathcal{L}_{\text{contrastive}} \quad (3)$$

while for ETTIN, it is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{contrastive}} \quad (4)$$

The choice of two losses allows us to better understand the impact of training on downstream retrieval performance.

3.4 Training Configuration

As the focus of our experiments is on understanding the behavior of embedding and model scaling, we use a fixed number of training epochs for all models in the same model family. For the BERT models, we use 50 epochs, while for the ETTIN models we use 15 as they have more training data per epoch. We use PyTorch [23] and Huggingface [33] for our implementation. We use AdamW as our optimizer with a learning rate of $2e-5$, and a linear learning rate scheduler with a warm-up ratio of 0.05. We use an effective batch size of 128 for BERT training and an effective batch size of 64 for ETTIN training. For all other hyperparameters, we use the default parameters from the Huggingface trainer.

A crucial choice for our analysis is which set of final embedding dimensions m to select. We decided to use multiples of each encoder’s native hidden size; for all models we used the multiples $\Phi : \{\frac{1}{4}, \frac{1}{2}, 1, 2, 4, 8, 16\}$ so that for a model with hidden size d the set of embedding dimensions is $\{\phi \cdot d : \phi \in \Phi\}$. For some of the larger BERT models, we included some additional lower multiples $\{\frac{1}{16}, \frac{1}{8}\}$ to understand how they behaved with additional compression.

3.5 Evaluation Data

Following prior work on scaling laws, we use evaluation data that is closely aligned with the training data [6, 8, 14]. Specifically, we use a combined version of TREC Deep Learning (DL) 2019, 2020, and 2021 [1–3] and MSMARCO Dev [22] for our main analysis. Both datasets feature an ad-hoc retrieval task with natural language queries, similar to the training data. MSMARCO Dev contains approximately 7k queries with incomplete labels that include one to two relevant passage per query. Our combined TREC DL contains 150 queries with far more complete labels and graded relevance labels, constructed through pooling.

We use these two datasets for deriving scaling laws and our core analysis as using less aligned data adds additional confounding

factors. With that said, understanding how performance on out-of-domain (i.e. unaligned) tasks is impacted by scaling is an important consideration. Thus, we include two additional evaluation tasks from the CRUMB benchmark [16, 29, 38] which have markedly different characteristics from the training data. The first task is Legal QA which has reasoning-intense legal questions as queries and legal statutes as documents. The second is Paper Retrieval where the queries are paragraphs describing features of a desired paper and the documents are paper titles and abstracts. We selected these tasks because the underlying tasks are still similar to the training task (i.e. retrieval), but the domain and query format are both different which should capture the models ability to generalize.

For TREC DL, MSMARCO Dev, and Legal QA we consider all passages in the QRELS with relevance above zero to be positive. For Paper Retrieval, we use the relevant passages provided by the binary QRELS.

3.6 Evaluation Metrics

3.6.1 Metrics. Following prior work investigating dense retrieval scaling laws [6], our primary evaluation metric is contrastive entropy. Contrastive entropy is a smooth measure of a retrieval model’s ability to separate a relevant document from irrelevant documents. The implementation is nearly identical to the contrastive learning loss we use:

$$L = -\log \frac{\exp(e_q \cdot e_{p^+})}{\exp(e_q \cdot e_{p^+}) + \sum_{p^- \in \mathcal{N}} \exp(e_q \cdot e_{p^-})}$$

Where e_q , e_{p^+} , and e_{p^-} represent the encoded query, positive passage, and negative passage, respectively. For the negatives, \mathcal{N} , following Fang et al. [6] we use 256 passages sampled uniformly at random from the entire corpus excluding known positive passages. For datasets where queries have multiple positives, we calculate contrastive entropy for each query-positive pair. The final contrastive entropy was found by taking the mean for all positives for the same query and then the mean over these query values. Like the contrastive loss we use during training, for ETTIN models we apply a temperature of $\tau = 0.02$.

To select the best checkpoint, we follow the approach used by Fang et al. [6] and run contrastive entropy, using our evaluation datasets, on each model checkpoint and select the best value. As we are interested in the scaling behavior this is a valid approach to remove potential noise.

To verify that the observed scaling behavior meaningfully reflects retrieval performance under standard ranking metrics, we study the empirical scaling results of a subset of our evaluation configuration based on reciprocal rank (RR@10) and recall (R@1000) in section 4.4.

3.6.2 Why Contrastive Entropy? As we aim to study embedding dimension scaling and derive scaling laws based on empirical results, our evaluation metric needs to be continuous and able to measure minor retrieval performance improvements. Traditional ranking metrics are not well-suited for this purpose as they are inherently discrete and insensitive to improvements beyond their cutoff parameter, while sensitive to small changes in ranking order. As contrastive entropy provides a continuous measure of a model’s ability to separate relevant from non-relevant documents and has

been shown in prior work [6] to exhibit strong correlation with conventional retrieval metrics, we select it as our primary evaluation measure.

Table 2: Joint scaling law parameters and goodness-of-fit (R^2) for BERT and Ettin model families on TREC DL combined and MSMARCO Dev.

Model	Set	A	B	α	β	δ	R^2
BERT	MS	114.887	0.800	1.887	1.247	0.013	0.975
	TREC	85.543	2.588	1.316	0.961	0.295	0.986
Ettin	MS	3.135	1.266	1.015	0.805	0.017	0.997
	TREC	2.010	2.444	0.505	0.588	0.450	0.978

4 Results and Analysis

In this section, we present the results of our experiments and derive scaling laws based on our empirical results.

4.1 Embedding Dimension Scaling Behavior

Given a model, we want to understand how scaling the embedding dimension impacts contrastive entropy, and, by extension, retrieval metrics. We do this by scaling the embedding dimension over a fixed range.

The results of this process are shown in Figure 1 for a selected BERT and Ettin model on our core evaluation tasks. The points, which represent the empirical results, follow a clear pattern; as embedding dimension increases, there is a decrease in contrastive entropy. Note that lower contrastive entropy is better and correlates to higher standard retrieval metrics. Although there is a clear decrease, as the embedding dimension continues to scale, the decrease begins to plateau. These results are interesting from a theoretical point of view as many theoretical limitations of embedding capacity are tied to the embedding dimension [17, 31] and one might expect that the gains in performance would continue with less saturation. From another theoretical perspective, the number of near orthogonal vectors increases exponentially with the number of dimensions [13], as near orthogonal vectors allow for more information with minimal interference, they are likely a good proxy for capacity. At some point, the amount of information necessary to perform a retrieval task is achieved and any further gains are diminishing. Additionally, as text and likely the concepts and features within text, follow a Zipfian distribution [39] at a certain point the vast majority of likely features are covered, and adding dimensions only helps for the infrequent long-tail features that have minimal impact on average evaluation performance.

The empirical results on their own are useful to better understand the impact of embedding dimension on retrieval performance, but in order to have a general sense of how scaling embeddings impacts performance it is useful to have an equation which can predict contrastive entropy given a model’s embedding dimension. Inspired by classical risk decomposition, we find that the contrastive entropy parameterized by the embedding dimension is well approximated by:

$$L(D) = \frac{A}{D^\alpha} + \delta_D \quad (5)$$

Where $L(D)$ is the predicted contrastive entropy and D is the embedding dimension. The parameters δ_D , α , and A are learned from the empirical data. From a risk decomposition perspective, δ_D represents the irreducible error of the task which is inherent due to ambiguity in queries and potential missing or incorrect labels. Additionally, as we only include a term to represent the impact of the embedding dimension the irreducible error will also account for the model’s limited ability to perfectly represent the task. We fit the empirical results to this equation using Scipy’s `curve_fit` which does a non-linear least squares to fit the function.

The result of fitting Equation (5) to our empirical results is shown in Figure 1 with the dashed lines. The fitted parameters and R^2 value, which is a statistical measure of the proportion of the variance in a dependent variable that is explained by the independent variable, for all models on MSMARCO Dev and TREC DL Combined can be observed in Table 1. The high values of R^2 mean that the variance in contrastive entropy is well explained by the embedding dimension; in other words, the predicted lines fit well. Nearly all models have $R^2 > 0.95$, but the largest Ettin on MSMARCO is quite a bit lower. This is likely because MSMARCO has lower contrastive entropy and so does the largest Ettin due to its size, this means noise is more likely to impact it and thus increase R^2 . A similar pattern, though less pronounced, is also true for the other large Ettin models on both MSMARCO and TREC DL. Overall, the fit is remarkably good across model families, model sizes, and evaluation data. The empirical adherence to the derived scaling law suggests that the embedding dimension scaling behavior follows a power law. This finding is in line with previous work on the scaling behavior of neural networks [6, 8, 14], which also follow power laws. **Knowing that performance gains from embedding dimension scaling follows a precise power law provides a powerful tool for architects of future dense retrieval models which enables them to predict downstream performance given various choices for embedding dimension.** Additionally, there are applications to prunable embeddings such as Matryoshka representations [18], which allow embeddings to be pruned to any size after encoding. With our scaling laws, the embedding size to use at retrieval time could be dynamically adjusted with a clear understanding of how the choice will impact retrieval performance.

4.2 Unified Model Parameter and Embedding Dimension Scaling Behavior

The prior scaling law assumes that we have a fixed model, but as shown by prior work on the scaling behavior of dense retrieval models [6], the model size has a substantial impact on retrieval performance. With this in mind, we want to understand the interplay between these two parameters, which is why we use a range of model sizes within each model family while conducting our experiments. The results on the combined TREC DL and MSMARCO Dev evaluation sets for the BERT and Ettin model families can be seen in Figures 2 and 3, respectively. The points represent the empirical results across model sizes as the embedding dimension is scaled. We observe that for similar embedding dimensions different model sizes have a significant impact on performance, although like scaling the embedding dimension the gains are diminishing. In this regard, our findings align with prior work on scaling dense




Figure 4: Out-of-domain evaluation on Paper Retrieval and Legal QA for BERT and Etnn families.

retrieval model sizes while keeping the embedding dimension fixed [6], which found the increase in performance from scaling model size follows a power law.

An interesting result of the joint embedding dimension and model size plot is that the scaling pattern remains largely consistent across different model sizes. For example, across all model sizes in Figure 2 we find that there is a sharp decrease from dimension 32 to 256 while beyond 1024 there is only a minimal decrease. This indicates that regardless of model scale the required embedding dimension to do a task well is similar and may be more related to the evaluation task than the model’s hidden size. This interpretation aligns with the prior discussion about why there is only limited improvement as embedding dimension increases. Specifically, for a given task there is some amount of information that is required to do well, but at a certain point the common features are accounted for and additional embedding dimensions only helps for infrequent features which, by definition, do not appear in many evaluation tasks.

Another takeaway is that it is very difficult to make up for limited model size with embedding dimension and vice versa. Once the embedding dimension is beyond a critical point where improvement begins to plateau, trying to match a larger model’s performance by scaling the embedding dimension becomes futile. For example in the TREC DL portion of Figure 2, the BERT-L4-H512 model can scale to an embedding dimension of 8k, but is still worse than the larger BERT-L8-H512 with an embedding dimension of 256. The same is also true for embedding dimension, but mostly at smaller embedding dimensions. For instance, when BERT-L8-H512 has an embedding dimension of 32 even the smallest BERT model can get lower performance with a larger embedding dimension, while this is not true when BERT-L8-H512 has an embedding dimension larger than 32.

Given the empirical results, as in the dimension only case, it is useful to have a predictable equation for contrastive entropy given both model size and embedding dimension. We do so by adding another term to the existing equation which accounts for

the improved performance as model size increases:

$$L(D, N) = \frac{A}{D^\alpha} + \frac{B}{N^\beta} + \delta \quad (6)$$

like Equation (5) D represents the embedding dimension, while N represents the model size based on the total number of parameters. The values of A , B , α , β , and δ are learned parameters. As before, δ represents the irreducible error of the task. We use the same approach to fit this scaling law as the dimension only scaling law. When fitting the scaling law we divide the model parameters by one million (e.g. 1,000,000 become 1.0) to reduce the scale of the parameters.

Applying the joint scaling law to our empirical results for the BERT model family on the combined TREC DL dataset and MS-MARCO Dev can be seen in Figure 2 with the dashed lines. Similarly, the results for the Etnn model family are shown in Figure 3.

One useful insight from our joint scaling laws is that **in many cases it is better to use a larger model with a smaller embedding dimension than a smaller model with a larger embedding dimension**. Although the exact trade-off is task specific, this finding is repeated across the models and evaluation tasks. A meaningful application of this is that embedding dimension is often the largest controllable factor in the size of the retrieval index and using smaller embeddings with a larger model could keep the same quality while halving the index size. We formalize a related idea in Section 4.5 where we find the optimal model size and embedding dimension for a fixed inference FLOPs budget.

4.3 Scaling Behavior for Aligned and Unaligned Evaluation Tasks

In the results shown so far, we focus on the contrastive entropy of “aligned” evaluation tasks. These are tasks that align with how the model was trained, including the training data and training objective. This practice is in line with prior work [6, 8, 14], but understanding how scaling impacts out-of-domain evaluation tasks is of significant importance for building generalizable retrieval

systems. In this section, we explore the scaling behaviors on out-of-domain evaluation tasks.

The results of evaluating both the ETTN and BERT model families on our two out-of-domain datasets Legal QA and Paper Retrieval are presented in Figure 4. The results show that generally as embedding dimension increases there is a decrease in contrastive entropy which matches the aligned results. However, unlike the aligned results the behavior is more erratic without an obvious, predictable pattern. Another unique characteristic is that contrastive entropy goes up when embedding dimension is increased beyond a certain dimension for several of the models. This behavior is especially frequent for the BERT model on the Paper Retrieval task. One explanation is that the Paper Retrieval task is more out-of-distribution for the BERT models which are only trained on short natural language questions, but this would not explain why the increase only happens as embedding dimensions become larger. Another explanation that better accounts for this uptick is that as the embedding dimension increases, a more complex similarity function can be learned. As the BERT models are trained with knowledge distillation, the larger embedding dimensions may become more aligned with the teacher and less able to generalize.

To investigate this hypothesis, we trained a set of BERT-L2-H128-A2 with only contrastive learning (CL) loss and no knowledge distillation. The contrastive entropy of this new CL only model compared with the original model can be viewed in Figure 5. We can see that the CL only version has nearly no uptick as embedding dimension scales, while the knowledge distillation variant does. This finding supports our hypothesis that the distillation loss enables the model learn a more specialized relevance function that is better suited for the aligned data, but which does not generalize as well. Crucially, the “overfitting” does not occur until a certain embedding dimension likely because prior to that point the representations lack the capacity to properly model the more specialized relevance function. This hypothesis is further supported by the absence of an uptick in the ETTN models which do not use knowledge distillation and prior work by Zeng et al. [37] that shows knowledge distillation can hurt generalization.

Even with the upticks in the BERT models the general trend remains the same: performance improves as we scale dimensions. **This means the scaling laws likely still broadly apply out-of-domain, just not as predictably.** These results mean scaling is likely a robust way to improve dense retrieval performance in a generalized way.

4.4 Scaling Behavior Based on Ranking Metrics

While contrastive entropy serves as our primary evaluation metric, it is important to assess whether the observed scaling trends translate to improvements in a full retrieval setting. Prior work on scaling behavior of dense retrieval models [6] has shown a strong positive correlation between contrastive entropy and standard ranking metrics. To verify that this correlation holds in our experimental setup, in this section we run an auxiliary retrieval evaluation and analyze scaling behavior based on classic retrieval metrics. Since we use contrastive entropy as a proxy for retrieval performance and exhaustive retrieval evaluation across all models and datasets would fall outside the scope of our study, we restrict this analysis to

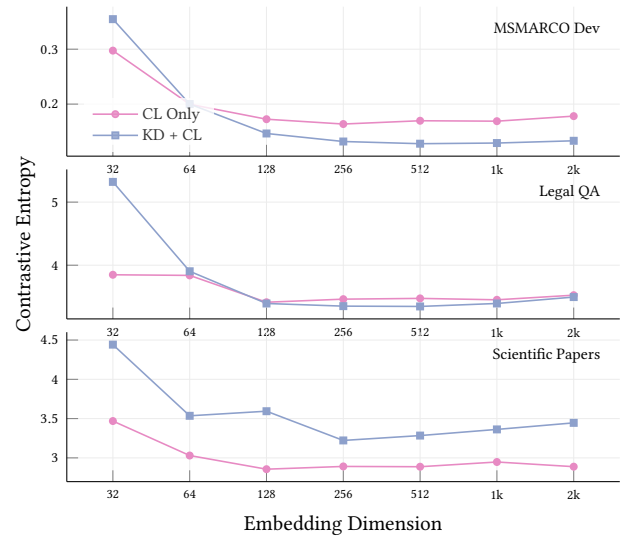


Figure 5: Contrastive entropy of BERT-L2-H128-A2 trained with only contrastive learning loss (CL Only) or contrastive learning loss and knowledge distillation loss (KD + CL) on MSMARCO Dev, Legal QA, and Scientific Papers.




Figure 6: Scaling behavior based on ranking metrics (RR@10 and R@1000) for BERT model family on MSMARCO Dev.

a representative subset of our experimental configurations. Specifically, we evaluate the BERT model family on the MSMARCO Dev dataset and report retrieval effectiveness in terms of reciprocal rank (RR@10) and recall (R@1000) as the embedding dimension is scaled. The results of this evaluation are presented in Figure 6. We can observe that the trend of performance improvement holds as we scale the embedding dimension, though the scaling behavior follows an increasing trend since we switched from contrastive entropy to ranking metrics. As expected from using discrete measures, the scaling behavior is not as smooth as depicted by contrastive entropy. This can be explained by the difference between the measurement objective of contrastive entropy and metrics such as R@1000 and RR@10. While ranking metrics evaluate the order of retrieved documents, they do not consider performance improvement beyond their cutoff point (10 for reciprocal rank and 1000 for recall in our




Figure 7: Optimal (\hat{N}, \hat{D}) under various FLOPs budgets and corpus sizes M . Across corpus sizes as FLOPs budget increases, both embedding size and model size increase, suggesting both are required for optimal performance.

setting), and they do not pick up on subtle improvements on document scores that do not alter the ranking order. Interestingly, this effect is especially evident when comparing the results in terms of RR@10 and R@1000. Reciprocal rank is inherently more restrictive in measuring retrieval performance here as it depends solely on the rank of the first relevant document and is capped at a cutoff of 10. In contrast, R@1000 accounts for all relevant documents retrieved within a substantially larger cutoff, allowing it to capture performance gains across a wider portion of the ranking. Consequently, the scaling trend with embedding dimension is more clearly reflected in R@1000 than in RR@10.

4.5 Cost Aware Embedding Dimension Selection

A useful application of our joint scaling law is that we can find the optimal trade-off between model size and embedding dimension at various compute budgets. In this section, we study the end-to-end trade-off between model capacity and retrieval cost under a fixed per-query FLOPs budget. Concretely, for each query we (i) run the encoder to produce a query embedding and (ii) score that embedding against a corpus of M precomputed document vectors. Let B denote the total FLOPs budget per query, T the query sequence length (tokens), N the number of model parameters, and D the embedding dimension. Following standard inference-cost approximations, we model the query-encoding cost as $C_{\text{enc}}(N, T) \approx 2NT$ FLOPs, and scoring as $C_{\text{score}}(M, D)$ FLOPs (counting one multiply-add per dimension per scored document). We enforce the compute budget $C_{\text{enc}} + C_{\text{score}} \leq B$ by sweeping an allocation parameter $\gamma \in (0, 1)$ that splits compute between encoding and scoring: $C_{\text{enc}} = \gamma B$ and $C_{\text{score}} = (1 - \gamma)B$, and then selecting the configuration that minimizes the predicted contrastive entropy under our fitted joint scaling law $\hat{L}(D, N)$:

$$\hat{\gamma} \in \arg \min_{\gamma \in (0, 1)} \hat{L}(D(\gamma), N(\gamma)), \quad (\hat{N}, \hat{D}) = (N(\hat{\gamma}), D(\hat{\gamma})).$$

Exhaustive scoring. For exhaustive (brute force) scoring over the full corpus, we use $C_{\text{score}}(M, D) \approx 2MD$. Under the γ -split, this

yields

$$N(\gamma) = \frac{\gamma B}{2T} \quad \text{and} \quad D(\gamma) = \frac{(1 - \gamma)B}{2M}.$$

ANN-style scoring. If approximate nearest neighbor (ANN) search is used, M can be replaced by an index-dependent proxy for the expected number of inner products evaluated per query. As a simple illustrative approximation, we can model the scoring cost as

$$C_{\text{score}}(M, D) \approx 2D \log(M),$$

which corresponds to evaluating a logarithmic number of candidate comparisons in M (up to method- and implementation-specific constants). Under the γ -split, this gives

$$N(\gamma) = \frac{\gamma B}{2T} \quad \text{and} \quad D(\gamma) = \frac{(1 - \gamma)B}{2 \log(M)}.$$

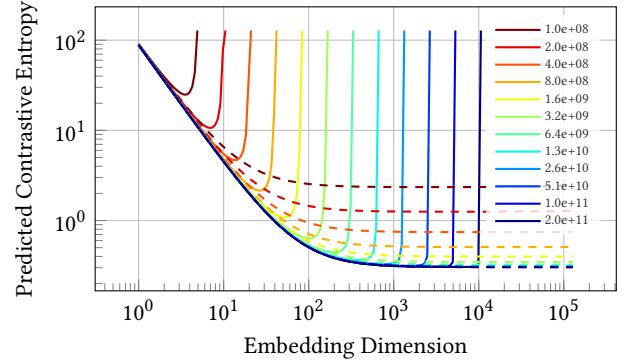


Figure 8: Predicted contrastive entropy using the joint scaling laws from Equation 6 for various embedding dimensions and FLOPs budgets. Solid lines represent exact scoring, dashed lines represent approximate-nearest-neighbor scoring. We use $T = 32$ and $M = 10^7$.

To keep this analysis tractable, we make some simplifying assumptions. First, we do not account for the projection to the final embedding dimension D in the encoding FLOPs C_{enc} . With that said, for most optimal setups the embedding dimension projection would be a very minor part of the overall encoding. We also assume ANN scoring would produce similar retrieval results to the exhaustive scoring. Although this is generally never completely true, in many use cases approximate approaches can be highly competitive with exhaustive approaches [30].

The results of the optimal embedding size and model size, with $M = \{10^5, 10^6, 10^7\}$ and $T = 32$, can be seen in Figure 7. We use the joint scaling law parameters fit on BERT using TREC DL Combined. The results show that across corpus sizes there is a clear pattern that for optimal FLOP allocation both model size and embedding dimension should be increased jointly. Intuitively, this makes sense as both axes, embedding dimension and model size, have diminishing returns as they are scaled so it is sensible the optimal approach scales both proportionally. We also see that for larger collections (i.e. larger values of M) the optimal embedding is smaller to account for the additional cost of scoring more documents.

Figure 8 shows the impact of varying the embedding dimension on the predicted contrastive entropy for different inference FLOP

budgets. We use the joint scaling law parameters fit on BERT using TREC DL Combined. **We observe that at low embedding dimensions, contrastive entropy is high because the embedding dimension is a bottleneck. As embedding dimension decreases so does the contrastive entropy until an inflection point where the contrastive entropy rapidly increases. This inflection point represent the optimal embedding dimension for a given FLOP budget.** After that point, the scoring becomes too expensive due to the embedding size and takes away from the FLOPs budget allocated to model size, which results in the sudden increase in contrastive entropy. The dashed lines represent the same setup as the solid lines but with approximate nearest neighbor scoring which we assume requires only $\log(M)$ inner-products. In the approximate setup, we can see that the optimal embedding dimension is far higher than with exhaustive retrieval, in fact the inflection point is not even visible on the graph. This means that with ANN scoring it is possible to substantially improve performance with the same FLOPs budget, this is especially true when using smaller budgets. **This insight might make a more efficient, but less effective, ANN approach paired with a larger embedding dimension preferable to the more conventional approach of using an embedding dimension inherited from the hidden state size with a high quality ANN approach.**

Our preliminary analysis is only the first step in applying the joint scaling laws we derived to inform dense retrieval modeling choices. Even with just this preliminary analysis, we have found interesting patterns and design guidelines for future developers of dense retrieval models. We believe that our scaling laws will be instrumental in scaling dense retrieval for the demanding new tasks of the future.

5 Conclusion

In this work, we presented a comprehensive empirical analysis of the relationship between embedding dimension and dense retrieval performance. By treating embedding dimension as an independent scaling axis rather than a static byproduct of the encoder architecture, we demonstrated that retrieval effectiveness follows a predictable power law as dimensionality increases. We validated this finding across two distinct model families, BERT and Ettn, and derived a joint scaling law that unifies model size and embedding dimension into a single predictive framework.

Our results offer critical insights for the design of efficient retrieval systems. We found that while increasing embedding dimension generally improves performance, it does so with diminishing returns. Crucially, our cost-aware analysis reveals that the standard practice of using the encoder’s native hidden size is rarely optimal. In many resource-constrained scenarios, utilizing a larger model with a projected, smaller embedding dimension yields superior retrieval quality compared to a smaller model with a larger embedding. Furthermore, we showed that when coupled with approximate nearest neighbor search, scaling embedding dimensions beyond the native hidden size can offer significant performance gains for a fixed compute budget.

While our scaling laws hold strongly for tasks aligned with the training distribution, our evaluation on out-of-domain collections

suggests that scaling behavior becomes less predictable when generalization is required, particularly for models trained with knowledge distillation. This highlights the importance of aligning model capacity with the complexity of the target retrieval task.

There are many avenues for future work in this area. A natural extension of this study is to investigate scaling laws for sparse embeddings, specifically analyzing the interplay between high-dimensional sparse representations and their sparsity ratios. Additionally, further research is needed to develop training objectives that encourage models to more uniformly utilize the geometric capacity of their embedding space, ensuring that every added dimension contributes meaningfully to retrieval effectiveness. We hope that the scaling laws and design guidelines presented in this work serve as a foundation for building the next generation of effective and efficient neural retrieval systems.

Acknowledgments

This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant #2402873, in part by the Office of Naval Research contract #N000142412612, and in part by the NSF Graduate Research Fellowships Program (GRFP) Award #1938059. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- [1] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020. Overview of the TREC 2020 deep learning track. In *Proceedings of the Twenty-Ninth Text REtrieval Conference (TREC 2020)*. NIST.
- [2] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020. Overview of the TREC 2020 Deep Learning Track. In *Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC 2020, Virtual Event [Gaithersburg, Maryland, USA], November 16–20, 2020 (NIST Special Publication, Vol. 1266)*, Ellen M. Voorhees and Angela Ellis (Eds.). National Institute of Standards and Technology (NIST). <https://trec.nist.gov/pubs/trec29/papers/OVERVIEW.DL.pdf>
- [3] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Jimmy Lin. 2021. Overview of the TREC 2021 Deep Learning Track. In *Proceedings of the Thirtieth Text REtrieval Conference, TREC 2021, online, November 15–19, 2021 (NIST Special Publication, Vol. 500–335)*, Ian Soboroff and Angela Ellis (Eds.). National Institute of Standards and Technology (NIST). <https://trec.nist.gov/pubs/trec30/papers/Overview-DL.pdf>
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 4171–4186. doi:10.18653/v1/N19-1423
- [5] Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The Faiss library. *CoRR abs/2401.08281* (2024). arXiv:2401.08281 doi:10.48550/ARXIV.2401.08281
- [6] Yan Fang, Jingtao Zhan, Qingyao Ai, Jiaxin Mao, Weihang Su, Jia Chen, and Yiqun Liu. 2024. Scaling Laws For Dense Retrieval. *CoRR abs/2403.18684* (2024). arXiv:2403.18684 doi:10.48550/ARXIV.2403.18684
- [7] Luyu Gao and Jamie Callan. 2021. Condenser: a Pre-training Architecture for Dense Retrieval. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 981–993. doi:10.18653/v1/2021.emnlp-main.75
- [8] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack W. Rae, and Laurent Sifre. 2022. Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NIPS '22)*. Article 2176, 15 pages.
- [9] Sebastian Hofstätter, Sophia Althammer, Michael Schröder, Mete Sertkan, and Allan Hanbury. 2020. Improving Efficient Neural Ranking Models

- with Cross-Architecture Knowledge Distillation. *CoRR* abs/2010.02666 (2020). arXiv:2010.02666 <https://arxiv.org/abs/2010.02666>
- [10] Sebastian Hofstätter, Sophia Althammer, Michael Schröder, Mete Sertkan, and Allan Hanbury. 2020. Improving Efficient Neural Ranking Models with Cross-Architecture Knowledge Distillation. *CoRR* abs/2010.02666 (2020). arXiv:2010.02666 <https://arxiv.org/abs/2010.02666>
- [11] 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* (2022). <https://openreview.net/forum?id=jKN1pXi7b0>
- [12] Hervé Jégou, Matthijs Douze, and Cordelia Schmid. 2011. Product Quantization for Nearest Neighbor Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 1 (2011), 117–128. doi:10.1109/TPAMI.2010.57
- [13] Paul C. Kainen and Věra Kůrková. 1993. Quasiorthogonal dimension of euclidean spaces. *Applied Mathematics Letters* 6, 3 (1993), 7–10. doi:10.1016/0893-9659(93)90023-G
- [14] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [15] 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)*. Association for Computational Linguistics, 6769–6781. doi:10.18653/v1/2020.emnlp-main.550
- [16] Julian Killingback and Hamed Zamani. 2025. Benchmarking Information Retrieval Models on Complex Retrieval Tasks. arXiv:2509.07253 [cs.LG] <https://arxiv.org/abs/2509.07253>
- [17] Julian Killingback, Hansi Zeng, and Hamed Zamani. 2025. Hypencoder: Hypernetworks for Information Retrieval. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*. Association for Computing Machinery, New York, NY, USA, 2372–2383. doi:10.1145/3726302.3729983
- [18] Aditya Kusupati, Gantavya Bhatt, Aniket Rege, Matthew Wallingford, Aditya Sinha, Vivek Ramanujan, William Howard-Snyder, Kaifeng Chen, Sham Kakade, Prateek Jain, et al. 2022. Matryoshka Representation Learning. In *Advances in Neural Information Processing Systems*.
- [19] Carlos Lassance, Hervé Déjean, Thibault Formal, and Stéphane Clinchant. 2024. SPLADE-v3: New baselines for SPLADE. *CoRR* abs/2403.06789 (2024). arXiv:2403.06789 doi:10.48550/ARXIV.2403.06789
- [20] Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. Sparse, Dense, and Attentional Representations for Text Retrieval. *Transactions of the Association for Computational Linguistics* 9 (2021), 329–345. doi:10.1162/TACL_A_00369
- [21] Aditya Menon, Sadeep Jayasumana, Ankit Singh Rawat, Seungyeon Kim, Sashank Reddi, and Sanjiv Kumar. 2022. In defense of dual-encoders for neural ranking. In *Proceedings of the 39th International Conference on Machine Learning*. PMLR. <https://proceedings.mlr.press/v162/menon22a.html>
- [22] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. In *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches*, Vol. 1773. CEUR-WS.org. https://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.), 8024–8035. <https://proceedings.neurips.cc/paper/2019/hash/bd8ca288fee7f92f2bfa9f7012727740-Abstract.html>
- [24] Jacob Portes, Connor Jennings, Erica Ji, Yuen Sasha Doubov, and Michael Carbin. 2025. Retrieval Capabilities of Large Language Models Scale with Pretraining FLOPs. arXiv:2508.17400 [cs.LG] <https://arxiv.org/abs/2508.17400>
- [25] Prafull Prakash, Julian Killingback, and Hamed Zamani. 2021. Learning Robust Dense Retrieval Models from Incomplete Relevance Labels. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.). ACM, 1728–1732. doi:10.1145/3404835.3463106
- [26] Ruiyang Ren, Yingqi Qu, Jing Liu, Wayne Xin Zhao, QiaoQiao She, Hua Wu, Haifeng Wang, and Ji-Rong Wen. 2021. RocketQAV2: A Joint Training Method for Dense Passage Retrieval and Passage Re-ranking. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2825–2835. doi:10.18653/v1/2021.emnlp-main.224
- [27] Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, Marine Carpuat, Marie-Catherine de Marneffe, and Iván Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, 3715–3734. doi:10.18653/v1/2022.naacl-main.272
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.), 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [29] Jianyou Wang, Kaicheng Wang, Xiaoyue Wang, Prudhviraj Naidu, Leon Bergen, and Ramamohan Paturi. 2023. Scientific Document Retrieval using Multi-level Aspect-based Queries. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). http://papers.nips.cc/paper_files/paper/2023/hash/78f9c04bdcb06f1ada3902912d8b64ba-Abstract-Datasets_and_Benchmarks.html
- [30] Mengzhao Wang, Xiaoliang Xu, Qiang Yue, and Yuxiang Wang. 2021. A comprehensive survey and experimental comparison of graph-based approximate nearest neighbor search. *Proc. VLDB Endow.* 14, 11 (July 2021), 1964–1978. doi:10.14778/3476249.3476255
- [31] Orion Weller, Michael Boratko, Iftekhar Naim, and Jinhyuk Lee. 2025. On the Theoretical Limitations of Embedding-based Retrieval. *CoRR* abs/2508.21038 (2025). arXiv:2508.21038 doi:10.48550/arXiv.2508.21038
- [32] Orion Weller, Kathryn Ricci, Marc Marone, Antoine Chaffin, Dawn J. Lawrie, and Benjamin Van Durme. 2025. Seq vs Seq: An Open Suite of Paired Encoders and Decoders. *CoRR* abs/2507.11412 (2025). arXiv:2507.11412 doi:10.48550/ARXIV.2507.11412
- [33] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2020. HuggingFace's Transformers: State-of-the-art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. doi:10.18653/v1/2020.emnlp-demos.6
- [34] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- [35] Xueguang Yang, Ronak Zhong, Tianqi Fang, and Jimmy Lin. 2020. Simple and Effective Unsupervised Redundancy Elimination to Compress Dense Vectors for Passage Retrieval. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 3154–3159. <https://aclanthology.org/2020.emnlp-main.255>
- [36] Hamed Zamani, Mostafa Dehghani, W. Bruce Croft, Erik Learned-Miller, and Jaap Kamps. 2018. From Neural Re-Ranking to Neural Ranking: Learning a Sparse Representation for Inverted Indexing. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (Torino, Italy) (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 497–506. doi:10.1145/3269206.3271800
- [37] Hansi Zeng, Julian Killingback, and Hamed Zamani. 2025. Scaling Sparse and Dense Retrieval in Decoder-Only LLMs. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (Padua, Italy) (SIGIR '25)*. Association for Computing Machinery, New York, NY, USA, 2679–2684. doi:10.1145/3726302.3730225
- [38] Lucia Zheng, Neel Guha, Javokhir Arifov, Sarah Zhang, Michal Skreta, Christopher D. Manning, Peter Henderson, and Daniel E. Ho. 2025. A Reasoning-Focused Legal Retrieval Benchmark. In *Proceedings of the 2025 Symposium on Computer Science and Law, CSLAW 2025, Munich, Germany, March 25-27, 2025*. ACM, 169–193. doi:10.1145/3709025.3712219
- [39] George K. Zipf. 1949. *Human Behaviour and the Principle of Least Effort*. Addison-Wesley.