

---

# On the Downstream Performance of Compressed Word Embeddings

---

Avner May    Jian Zhang    Tri Dao    Christopher Ré  
Department of Computer Science, Stanford University  
{avnermay, zjian, trid, chrismre}@cs.stanford.edu

## Abstract

Compressing word embeddings is important for deploying NLP models in memory-constrained settings. However, understanding what makes compressed embeddings perform well on downstream tasks is challenging—existing measures of compression quality often fail to distinguish between embeddings that perform well and those that do not. We thus propose the *eigenspace overlap score* as a new measure. We relate the eigenspace overlap score to downstream performance by developing generalization bounds for the compressed embeddings in terms of this score, in the context of linear and logistic regression. We then show that we can lower bound the eigenspace overlap score for a simple uniform quantization compression method, helping to explain the strong empirical performance of this method. Finally, we show that by using the eigenspace overlap score as a selection criterion between embeddings drawn from a representative set we compressed, we can efficiently identify the better performing embedding with up to  $2\times$  lower selection error rates than the next best measure of compression quality, and avoid the cost of training a model for each task of interest.

## 1 Introduction

In recent years, *word embeddings* [22, 28, 23, 29, 10] have brought large improvements to a wide range of applications in natural language processing (NLP) [1, 5, 37]. However, these word embeddings can occupy a large amount of memory, making it expensive to deploy them in data centers, and impractical to use them in memory-constrained environments like smartphones. To reduce and amortize these costs, embeddings can be compressed [e.g., 33] and shared across many downstream tasks [7]. Recently, there have been numerous successful methods proposed for compressing embeddings; these methods take a variety of approaches, ranging from compression using k-means clustering [2] to dictionary learning using neural networks [33, 6].

The goal of this work is to gain a deeper understanding of what makes compressed embeddings perform well on downstream tasks. Practically, this understanding could allow for evaluating the quality of a compressed embedding without having to train a model for each task of interest. Our work is motivated by two surprising empirical observations: First, we find that existing ways [40, 3, 41] of measuring the quality of compressed embeddings do not effectively explain the relative *downstream* performance of different compressed embeddings—for example, failing to discriminate between embeddings that perform well and those that do not. Second, we observe that a simple uniform quantization method can match or outperform the state-of-the-art deep compositional code learning method [33] and the k-means compression method [2] in terms of downstream performance. These observations suggest that there is currently an incomplete understanding of what makes a compressed embedding perform well on downstream tasks. One way to narrow this gap in our understanding is to find a measure of compression quality that (i) is directly related to generalization performance, and (ii) can be used to analyze the performance of uniformly quantized embeddings.

Here we introduce the *eigenspace overlap score* as a new measure of compression quality, and show that it satisfies the above two desired properties. This score measures the degree of overlap between the subspaces spanned by the eigenvectors of the Gram matrices of the compressed and uncompressed embedding matrices. Our theoretical contributions are two-fold, addressing the surprising observations and desired properties discussed above: First, we prove generalization bounds for the compressed embeddings in terms of the eigenspace overlap score in the context of linear and logistic regression, revealing a direct connection between this score and downstream performance. Second, we prove that in expectation uniformly quantized embeddings attain a high eigenspace overlap score with the uncompressed embeddings at relatively high compression rates, helping to explain their strong performance. Inspired by these theoretical connections between the eigenspace overlap score and generalization performance, we propose using this score as a selection criterion for efficiently picking among a set of compressed embeddings, without having to train a model for each task of interest using each embedding.

We empirically validate our theoretical contributions and the efficacy of our proposed selection criterion by showing three main experimental results: First, we show the eigenspace overlap score is more predictive of downstream performance than existing measures of compression quality [40, 3, 41]. Second, we show uniform quantization consistently matches or outperforms all the compression methods to which we compare [2, 33, 15], in terms of both the eigenspace overlap score and downstream performance. Third, we show the eigenspace overlap score is a more accurate criterion for choosing between compressed embeddings than existing measures; specifically, we show that when choosing between embeddings drawn from a representative set we compressed [2, 33, 11, 15], the eigenspace overlap score is able to identify the one that attains better downstream performance with up to  $2\times$  lower selection error rates than the next best measure of compression quality. We consider several baseline measures of compression quality: the Pairwise Inner Product (PIP) loss [40], and two spectral measures of approximation error between the embedding Gram matrices [3, 41]. Our results are consistent across a range of NLP tasks [32, 18, 37], embedding types [28, 23, 10], and compression methods [2, 33, 11].

The rest of this paper is organized as follows. In Section 2 we review background on word embedding compression methods and existing measures of compression quality, and present the two motivating empirical observations. In Section 3 we present the eigenspace overlap score along with our corresponding theoretical contributions, and propose to use the eigenspace overlap score as a selection criterion. In Section 4, we show the results from our extensive experiments validating the practical significance of our theoretical contributions, and the efficacy of our proposed selection criterion. We present related work in Section 5, and conclude in Section 6.

## 2 Background and Motivation

We first review different compression methods in Section 2.1 and existing ways to measure the quality of a compressed embedding relative to the uncompressed embedding in Section 2.2. We then show in Section 2.3 that existing measures of compression quality do not satisfactorily explain the relative downstream performance of existing compression methods; this motivates our work to better understand the downstream performance of compressed embeddings.

### 2.1 Embedding Compression Methods

We now discuss a number of compression methods for word embeddings. For the purposes of this paper, the goal of an embedding compression method  $C(\cdot)$  is to take as input an uncompressed embedding  $X \in \mathbb{R}^{n \times d}$ , and produce as output a compressed embedding  $\tilde{X} := C(X) \in \mathbb{R}^{n \times k}$  which uses less memory than  $X$ , but attains similar performance to  $X$  when used in downstream models. Here,  $n$  denotes the vocabulary size,  $d$  and  $k$  the uncompressed and compressed dimensions.

**Deep Compositional Code Learning (DCCL)** The DCCL method [33] uses a dictionary learning approach to represent a large number of word vectors using a much smaller number of basis vectors. The dictionaries are trained using an autoencoder-style architecture to minimize the embedding matrix reconstruction error. A similar approach was independently proposed by Chen et al. [6].

**K-means Compression** The k-means algorithm can be used to compress word embeddings by first clustering all the scalar entries in the word embedding matrix, and then replacing each scalar with the closest centroid [2]. Using  $2^b$  centroids allows for storing each matrix entry using only  $b$  bits.

**Dimensionality Reduction** One can train an embedding with a lower dimension, or use a method like principal component analysis (PCA) to reduce the dimensionality of an existing embedding.

**Uniform Quantization** To compress real numbers, uniform quantization divides an interval into sub-intervals of equal size, and then (deterministically or stochastically) rounds the numbers in each sub-interval to one of the boundaries [11, 13]. To apply uniform quantization to embedding compression, we propose to first determine the optimal threshold at which to clip the extreme values in the word embedding matrix, and then uniformly quantize the clipped embeddings within the clipped interval. For more details about uniform quantization and how we use it to compress embeddings, see Appendices A.1 and D.3 respectively.

## 2.2 Measures of Compression Quality

We review ways of measuring the compression quality of a compressed embedding relative to the uncompressed embedding. For our purposes, an ideal measure would consider a compressed embedding to have high quality when it is likely to perform similarly to the uncompressed embedding on downstream tasks, and low quality otherwise. Such a measure would shed light on what determines the downstream performance of a compressed embedding, and give us a way of measuring the quality of a compressed embedding without having to train a downstream model for each task.

Several of the measures discussed below are based on comparing the pairwise inner product (Gram) matrices of the compressed and uncompressed embeddings. The Gram matrices of embeddings are natural to consider for two reasons: First, the loss function for training word embeddings typically only considers dot products between embedding vectors [22, 28]. Second, one can view word embedding training as implicit matrix factorization [20], and thus comparing the Gram matrices of two embedding matrices is similar to comparing the matrices these embeddings are implicitly factoring. We now review several existing ways of measuring compression quality.

**Word Embedding Reconstruction Error** The first and simplest way of comparing two embeddings  $X$  and  $\tilde{X}$  is to measure the reconstruction error  $\|X - \tilde{X}\|_F$ . Note that in order to be able to use this measure of quality,  $X$  and  $\tilde{X}$  must have the same dimension.

**Pairwise Inner Product (PIP) Loss** Given  $XX^T$  and  $\tilde{X}\tilde{X}^T$ , the Gram matrices of the uncompressed and compressed embeddings, the *Pairwise Inner Product (PIP) Loss* [40] is defined as  $\|XX^T - \tilde{X}\tilde{X}^T\|_F$ . This measure of quality was recently proposed to explain the existence of an optimal dimension for word embeddings, in terms of a bias-variance trade-off for the PIP loss.

**Spectral Approximation Error** A symmetric matrix  $A$  is defined [41] to be a  $(\Delta_1, \Delta_2)$ -spectral approximation of another symmetric matrix  $B$  if it satisfies  $(1 - \Delta_1)B \preceq A \preceq (1 + \Delta_2)B$  (in the semidefinite order). Zhang et al. [41] show that if  $\tilde{X}\tilde{X}^T + \lambda I$  is a  $(\Delta_1, \Delta_2)$ -spectral approximation of  $XX^T + \lambda I$  for sufficiently small values of  $\Delta_1$  and  $\Delta_2$ , then the linear model trained using  $\tilde{X}$  and regularization parameter  $\lambda$  will attain similar generalization performance to the model trained using  $X$ . Avron et al. [3] use a single scalar  $\Delta$  in place of  $\Delta_1$  and  $\Delta_2$ , and use this scalar as a measure of approximation error, while Zhang et al. [41] consider  $\Delta_1$  and  $\Delta_2$  independently, and use the quantity  $\Delta_{\max} := \max(\frac{1}{1-\Delta_1}, \Delta_2)$  to measure approximation error.

## 2.3 Two Motivating Empirical Observations

We now present two empirical observations which illustrate the need to better understand the downstream performance of models trained using compressed embeddings. In these experiments we compare the downstream performance of the methods introduced in Section 2.1, and attempt to use the measures of compression quality from Section 2.2 to explain the relative performance of these compression methods. Our observations reveal that explaining the downstream performance of compressed embeddings is challenging. We now provide an overview of these two observations; for a more thorough presentation of these results, see Section 4.

- First, we observe that the downstream performance of embeddings compressed using the various methods from Section 2.1 cannot be satisfactorily explained in terms of any of the existing measures of compression quality described in Section 2.2. For example, in Figure 1 we see that on GloVe embeddings [28], the uniform quantization method with compression rate  $32\times$  can have over  $1.3\times$  higher PIP loss than dimensionality reduction with compres-

sion rate  $6\times$ , while attaining better downstream performance by over 2.5 F1 points on the Stanford Question Answering Dataset (SQuAD) [32]. Furthermore, the PIP loss and the two spectral measures of approximation error  $\Delta$  and  $\Delta_{\max}$  only achieve Spearman correlation absolute values of 0.49, 0.46, and 0.62 with the question answering test F1 score, respectively (Table 1). These results show that existing measures of compression quality correlate relatively poorly with downstream performance.

- Our second observation is that the simple uniform quantization method matches or outperforms the more complex DCCL and k-means compression methods across a number of tasks, embedding types, and compression ratios. For example, with a compression ratio of  $32\times$ , uniform quantization attains an average F1 score 0.47 points below the uncompressed GloVe embeddings on the Stanford Question Answering Dataset [32], while the DCCL method [33] is 0.43 points below.




Figure 1: The PIP loss does not satisfactorily explain the relative downstream performance of different compression methods.

These two observations suggest the need to better understand the downstream performance of compressed embeddings. Toward this end, we focus on finding a measure of compression quality with the properties that (i) we can directly relate it to generalization performance, and (ii) we can use it to analyze the performance of uniformly quantized embeddings.

### 3 A New Measure of Compression Quality

To better understand what properties of compressed embeddings determine their downstream performance, and to help explain the motivating empirical observations above, we introduce the *eigenspace overlap score*, and show that it satisfies the two desired properties described above. In Section 3.1 we present generalization bounds for compressed embeddings in the context of linear and logistic regression, in terms of the eigenspace overlap score between the compressed and uncompressed embeddings. In Section 3.2 we show that in expectation, uniformly quantized embeddings attain high eigenspace overlap scores, helping to explain their strong downstream performance. Based on the connection between the eigenspace overlap score and downstream performance, in Section 3.3 we propose using this score as a way of efficiently selecting among different compressed embeddings.

#### 3.1 The Eigenspace Overlap Score and Generalization Performance

We begin by defining the eigenspace overlap score, which measures how well a compressed embedding approximates an uncompressed embedding. We then present our theoretical results relating the generalization performance of compressed embeddings to their eigenspace overlap scores.

##### 3.1.1 The Eigenspace Overlap Score

We now define the eigenspace overlap score, and discuss the intuition behind this definition.

**Definition 1.** Given two full-rank embedding matrices  $X \in \mathbb{R}^{n \times d}$ ,  $\tilde{X} \in \mathbb{R}^{n \times k}$ , whose Gram matrices have eigendecompositions  $XX^T = U\Lambda U^T$ ,  $\tilde{X}\tilde{X}^T = \tilde{U}\tilde{\Lambda}\tilde{U}^T$  for  $U \in \mathbb{R}^{n \times d}$ ,  $\tilde{U} \in \mathbb{R}^{n \times k}$ , we define the eigenspace overlap score  $\mathcal{E}(X, \tilde{X}) := \frac{1}{\max(d, k)} \|U^T \tilde{U}\|_F^2$ .

This score quantifies the similarity between the subspaces spanned by the eigenvectors with nonzero eigenvalues of  $\tilde{X}\tilde{X}^T$  and  $XX^T$ . In particular, assuming  $k \leq d$ , it measures the ratio between the squared Frobenius norm of  $U$  before and after being projected onto  $\tilde{U}$ . It attains a maximum value of one when  $\text{span}(U) = \text{span}(\tilde{U})$ , and a minimum value of zero when these two spans are orthogonal. Computing this score takes time  $O(n \max(d, k)^2)$ , as it requires computing the singular value decompositions (SVDs) of  $X$  and  $\tilde{X}$ . As is clear from the definition, the eigenspace overlap score only depends on the left singular vectors of the two embedding matrices. To better understand why this is a desirable property, consider two embedding matrices  $X$  and  $\tilde{X}$  with the same left singular vectors. It follows that the output of any linear model over  $X$  can be exactly matched by the output of a linear model over  $\tilde{X}$ ; if we consider the SVDs  $X = USV^T$ ,  $\tilde{X} := U\tilde{S}\tilde{V}^T$ , then for any

parameter vector  $w \in \mathbb{R}^d$  over  $X$ ,  $\tilde{w} := \tilde{V}\tilde{S}^{-1}SV^Tw$  gives  $Xw = \tilde{X}\tilde{w}$ . This observation shows how central the left singular vectors of an embedding matrix are to the set of models which use this matrix, and thus why it is reasonable for the eigenspace overlap score to only consider the left singular vectors. In Appendix B.3 we discuss this score's robustness to perturbations, while in Appendix B.4 we discuss the connection between this score and a variant of embedding reconstruction error.

### 3.1.2 Generalization Results

We now present our theoretical results relating the difference in generalization performance between models trained on compressed vs. uncompressed embeddings, in terms of the eigenspace overlap score. For these results, we consider an average-case analysis in the context of fixed design linear regression, for both the squared loss function and for any Lipschitz continuous loss function (e.g., logistic loss). We consider the fixed design setting for ease of analysis; for example, when using the squared loss there is a closed-form expression for a regressor's generalization performance. Before presenting our results in Theorems 1 and 2 for the two types of loss functions, we briefly review fixed design linear regression, and discuss the average-case setting we consider.

In fixed design linear regression, we observe a set of labeled points  $\{(x_i, y_i)\}_{i=1}^n$  where the observed labels  $y_i = \bar{y}_i + \epsilon_i \in \mathbb{R}$  are perturbed from the true labels  $\bar{y}_i$  with independent noise  $\epsilon_i$  with mean zero and variance  $\sigma^2$ . If we let  $x_i \in \mathbb{R}^d$  denote the  $i^{\text{th}}$  row of the matrix  $X \in \mathbb{R}^{n \times d}$  with SVD  $X = USV^T$ , let  $y$  and  $\bar{y}$  in  $\mathbb{R}^n$  denote the perturbed and true label vectors, and let  $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  be a convex loss function, we can define  $f_{X,\epsilon}$  as the linear model which minimizes the empirical loss:  $f_{X,\epsilon}(x) := x^T w^*$  where  $w^* := \arg \min_{w \in \mathbb{R}^d} \sum_{i=1}^n \ell(x_i^T w, y_i)$ . When the loss function is the squared loss, we can use the closed-form solution  $w^* = (X^T X)^{-1} X^T y$  to show that the expected loss of  $f_{X,\epsilon}$  is equal to  $\mathcal{R}_{\bar{y}}(X) := \mathbb{E}_{\epsilon} \left[ \frac{1}{n} \sum_{i=1}^n \ell(f_{X,\epsilon}(x_i), \bar{y}_i) \right] = \frac{1}{n} (\|\bar{y}\|^2 - \|U^T \bar{y}\|^2 + d\sigma^2)$ ; for the derivation, see Appendix A.2. If we instead consider any Lipschitz continuous convex loss function (e.g., the logistic loss<sup>1</sup>) there may not be a closed-form solution for the parameter vector  $w^*$ , but we can still derive upper bounds on the expected loss in this setting (see Theorem 2).

We consider average-case analysis for two reasons: First, in the setting where one would like to use the same compressed embedding across many tasks (i.e., different label vectors  $\bar{y}$ ), an average-case result describes the average performance across these tasks. Second, for both empirical and theoretical reasons we argue that worst-case bounds are too loose to explain our empirical observations. Empirically, we observe that compressed embeddings with large values of  $\Delta_1$  and  $\Delta_2$  (defined in Section 2.2) can still attain strong generalization performance (Appendix E.6), even though these values imply large worst-case bounds on the generalization error [41]. From a theoretical perspective, worst-case bounds must account for all possible label vectors, including those chosen adversarially. For example, if there exists a single direction in  $\text{span}(U)$  orthogonal to  $\text{span}(\tilde{U})$  (which always occurs when  $\dim(\tilde{U}) < \dim(U)$ ) the label vector  $\bar{y}$  can be in this direction, resulting in large generalization error for  $\tilde{X}$  and small generalization error for  $X$ . Thus, we consider an average-case analysis in which we assume  $\bar{y}$  is a random label vector in  $\text{span}(U)$ . We consider this setting because we are most interested in the situation where we know the uncompressed embedding matrix  $X$  performs well (in this case,  $\mathcal{R}_{\bar{y}}(X) = d\sigma^2/n$ ), and we would like to understand how well  $\tilde{X}$  can do.<sup>2</sup>

We now present our result for the squared loss. To maintain a constant signal ( $\bar{y}$ ) to noise ( $\epsilon$ ) ratio for different embedding matrix sizes, we define  $c \in \mathbb{R}$  as the scalar for which  $\sigma^2 = c^2 \cdot \mathbb{E}_{\bar{y}} \left[ \frac{1}{n} \sum_{i=1}^n \bar{y}_i^2 \right]$ . Thus, when  $c = 1$  the entries of the true label vector on average have the same variance as the noise.

**Theorem 1.** *Let  $X = USV^T \in \mathbb{R}^{n \times d}$  be the singular value decomposition of a full-rank embedding matrix  $X$ , and let  $\tilde{X} \in \mathbb{R}^{n \times k}$  be another full-rank embedding matrix. Let  $\bar{y} = Uz \in \mathbb{R}^n$  denote a random label vector in  $\text{span}(U)$ , where  $z$  is random with zero mean and identity covariance matrix. Letting  $\sigma^2 = c^2 \cdot \mathbb{E}_{\bar{y}} \left[ \frac{1}{n} \sum_{i=1}^n \bar{y}_i^2 \right] = c^2 \frac{d}{n}$  denote the variance of the label noise, it follows that*

$$\mathbb{E}_{\bar{y}} \left[ \mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X) \right] = \frac{d}{n} \cdot \left( 1 - \mathcal{E}(X, \tilde{X}) \right) - c^2 \cdot \frac{d(d-k)}{n^2}. \quad (1)$$

<sup>1</sup>We consider the logistic loss  $\ell(z', z) := -(\sigma(z) \log(\sigma(z')) + (1 - \sigma(z)) \log(1 - \sigma(z')))$ , where here  $\sigma: \mathbb{R} \rightarrow \mathbb{R}$  denotes the sigmoid function, and  $z$  and  $z'$  both represent logits. If  $z' := w^T x$  is bounded (which occurs when the weight vector and data are both bounded), this loss is Lipschitz continuous in both arguments.

<sup>2</sup>The difference between average-case and worst-case analysis is central to understanding the difference between  $(\Delta_1, \Delta_2)$ -spectral approximation (which yields worst-case generalization bounds) [41] and the eigenspace overlap score (which yields average-case generalization bounds).

This theorem reveals that a larger eigenspace overlap score  $\mathcal{E}(X, \tilde{X})$  results in better expected loss for the compressed embedding. Note that if we focus on the low-dimensional and low-noise setting, where  $d \ll n$  and  $c^2 = O(1)$ , we can effectively ignore the term  $c^2 \frac{d(d-k)}{n^2} = O(d^2/n^2)$ , and the generalization performance is determined by the eigenspace overlap score.

We now present a result analogous to Theorem 1 for Lipschitz continuous loss functions.

**Theorem 2.** *Let  $X \in \mathbb{R}^{n \times d}$ ,  $\tilde{X} \in \mathbb{R}^{n \times k}$ ,  $\bar{y} \in \mathbb{R}^n$ , and  $c \in \mathbb{R}$  be defined as in Theorem 1. Let  $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  be a convex non-negative loss function which is  $L$ -Lipschitz continuous in both arguments and satisfies  $\arg \min_v \ell(v', v) = v \ \forall v \in \mathbb{R}$ . It follows that*

$$\mathbb{E}_{\bar{y}} \left[ \mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X) \right] \leq \frac{L\sqrt{d}}{\sqrt{n}} \left( \sqrt{1 - \mathcal{E}(X, \tilde{X})} + 2c \right).$$

Similarly to Theorem 1, we see that a larger eigenspace overlap score results in a tighter bound on the generalization performance of the compressed embeddings. See Appendix B for the proofs for Theorems 1 and 2, where we consider the more general setting of  $z$  having arbitrary covariance.

### 3.2 The Eigenspace Overlap Score and Uniform Quantization

To help explain the strong downstream performance of uniformly quantized embeddings, in this section we present a lower bound on the expected eigenspace overlap score for uniformly quantized embeddings. Combining this result with Theorem 1 directly provides a guarantee on the performance of the uniformly quantized embeddings.

To prove this bound on the eigenspace overlap score, we use the Davis-Kahan  $\sin(\Theta)$  theorem [8], which upper bounds the amount the eigenvectors of a matrix can change after the matrix is perturbed, in terms of the perturbation magnitude. Because for uniform quantization we can exactly characterize the magnitude of the perturbation, this theorem allows us to bound the eigenspace overlap score of uniformly quantized embeddings. Note that we assume unbiased stochastic rounding is used for the uniform quantization (see [13] or Appendix A.1). We now present the result (proof in Appendix C):

**Theorem 3.** *Let  $X \in \mathbb{R}^{n \times d}$  be a bounded embedding matrix with  $X_{ij} \in [-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}]^3$  and smallest singular value  $\sigma_{\min} = a\sqrt{n/d}$ , for  $a \in (0, 1]$ .<sup>4</sup> Let  $\tilde{X}$  be an unbiased stochastic uniform quantization of  $X$ , where  $b$  bits are used per entry. Then for  $n \geq \max(33, d)$ , we can lower bound the expected eigenspace overlap score of  $\tilde{X}$ , over the randomness of the stochastic quantization, as follows:*

$$\mathbb{E} \left[ 1 - \mathcal{E}(X, \tilde{X}) \right] \leq \frac{20}{(2^b - 1)^2 a^4}.$$

A consequence of this theorem is that with only a logarithmic number of bits  $b \geq \log_2 \left( \frac{\sqrt{20}}{a^2 \sqrt{\epsilon}} + 1 \right)$ , uniform quantization can attain an expected eigenspace overlap score of at least  $1 - \epsilon$ . This helps explain the strong downstream performance of uniform quantization at high compression rates.

In Appendix C.2 we empirically validate that the scaling of the eigenspace overlap score with respect to the quantities in Theorem 3 matches the theory; we show  $1 - \mathcal{E}(X, \tilde{X})$  drops as the precision  $b$  and the scalar  $a$  are increased, and is relatively unaffected by changes to the vocabulary size  $n$  and dimension  $d$ .

### 3.3 The Eigenspace Overlap Score as a Selection Criterion

Due to the theoretical connections between generalization performance and the eigenspace overlap score, we propose using the eigenspace overlap score as a selection criterion between different compressed embeddings. Specifically, the algorithm we propose takes as input an uncompressed embedding along with two or more compressed versions of this embedding, and returns the compressed embedding with the highest eigenspace overlap score to the uncompressed embedding. Ideally, a selection criterion should be both accurate and robust. For each downstream task, we consider

<sup>3</sup>This bound on the entries of  $X$  results in the entries of its Gram matrix being bounded by a constant independent of  $d$ .

<sup>4</sup>The maximum possible value of  $\sigma_{\min}$  is  $\sqrt{n/d}$ , which occurs when  $\|X\|_F^2 = n$  and  $\sigma_{\min} = \sigma_{\max}$ .




Figure 2: **Downstream performance vs. measures of compression quality.** We plot the performance of compressed fastText embeddings on the SQuAD question answering task as a function of different measures of compression quality. The eigenspace overlap score  $\mathcal{E}$  demonstrates better alignment with downstream performance across compression methods than the other measures. We quantify the degree of alignment using the Spearman correlation  $\rho$ , and include  $\rho$  in the plot titles.

*accuracy* as the fraction of cases where a criterion selects the best-performing embedding on the task. We quantify the *robustness* as the maximum observed performance difference between the selected embedding and the one which performs the best on a downstream task. In Section 4.3, we empirically validate that the eigenspace overlap score is a more accurate and robust criterion than existing measures of compression quality.

## 4 Experiments

We empirically validate our theory relating the eigenspace overlap score with generalization performance, our analysis on the strong performance of uniform quantization, and the efficacy of the eigenspace overlap score as an embedding selection criterion. We first demonstrate that this score correlates better with downstream performance than existing measures of compression quality in Section 4.1. We then demonstrate in Section 4.2 that uniform quantization consistently matches or outperforms the compression methods to which we compare, both in terms of the eigenspace overlap score and downstream performance. In Section 4.3, we show that the eigenspace overlap score is a more accurate and robust selection criterion than other measures of compression quality.

**Experiment setup** We evaluate compressed versions of publicly available 300-dimensional fastText and GloVe embeddings on question answering and sentiment analysis tasks, and compressed 768-dimensional WordPiece embeddings from the pre-trained case-sensitive BERT<sub>BASE</sub> model [10] on tasks from the General Language Understanding Evaluation (GLUE) benchmark [37]. We use the four compression methods discussed in Section 2: DCCL, k-means, dimensionality reduction, and uniform quantization.<sup>5</sup> For the tasks, we consider question answering using the DrQA model [5] on the Stanford Question Answering Dataset (SQuAD) [32], sentiment analysis using a CNN model [18] on all the datasets used by Kim [18], and language understanding using the BERT<sub>BASE</sub> model on the tasks in the GLUE benchmark [37]. We present results on the SQuAD dataset, the largest sentiment analysis dataset (SST-1 [34]) and the two largest GLUE tasks (MNLI and QQP) in this section, and include the results on the other sentiment analysis and GLUE tasks in Appendix E. We evaluate downstream performance using the F1 score for question answering, accuracy for sentiment analysis, and the standard evaluation metric for each GLUE task (Table 5 in Appendix D). Across embedding types and tasks, we first compress the pre-trained embeddings, and then train the non-embedding model parameters in the standard manner for each task, keeping the embeddings fixed throughout training. For the GLUE tasks, we add a linear layer on top of the final layer of the pre-trained BERT model (as in [10]), and then fine-tune the non-embedding model parameters.<sup>6</sup> For more details on the various embeddings, tasks, and hyperparameters we use, see Appendix D.

### 4.1 The Eigenspace Overlap Score and Downstream Performance

To empirically validate the theoretical connection between the eigenspace overlap score and downstream performance, we show that the eigenspace overlap score correlates better with downstream performance than the existing measures of compression quality discussed in Section 2. Thus, even though our analysis is for linear and logistic regression, we see the eigenspace overlap score also has strong empirical correlation with downstream performance on tasks using neural network models.

<sup>5</sup>For dimensionality reduction, we use PCA for fastText and BERT embeddings (compression rates: 1, 2, 4, 8), and publicly available lower-dimensional embeddings for GloVe (compression rates: 1, 1.5, 3, 6).

<sup>6</sup>Freezing the WordPiece embeddings does not observably affect performance (see Appendix E.1).

Table 1: **Spearman correlation between measures of compression quality and downstream performance.** For each measure of compression quality, we show the absolute value of its Spearman correlation with downstream performance, on the SQuAD (question answering), SST-1 (sentiment analysis), MNLI (natural language inference), and QQP (question pair matching) tasks. We see that the eigenspace overlap score  $\mathcal{E}$  attains stronger correlation than the other measures.

Dataset	SQuAD		SST-1		MNLI	QQP
Embedding	GloVe	fastText	GloVe	fastText	BERT WordPiece	BERT WordPiece
PIP loss	0.49	0.34	0.46	0.25	0.45	0.45
$\Delta$	0.46	0.31	0.33	0.29	0.44	0.36
$\Delta_{\max}$	0.62	0.72	0.51	0.60	0.86	0.86
$1 - \mathcal{E}$	<b>0.81</b>	<b>0.91</b>	<b>0.75</b>	<b>0.73</b>	<b>0.92</b>	<b>0.93</b>

In Figure 2 we present results for question answering (SQuAD) performance for compressed fastText embeddings as a function of the various measures of compression quality. In each plot, for each combination of compression rate and compression method, we plot the average compression quality measure ( $x$ -axis) and the average downstream performance ( $y$ -axis) across the five random seeds used (error bars indicate standard deviations). If the ranking based on the measure of compression quality was identical to the ranking based on downstream performance, we would see a monotonically decreasing sequence of points. As we can see from the rightmost plot in Figure 2, the downstream performance decreases smoothly as the eigenspace overlap value decreases; the downstream performance does not align as well with the other measures of compression quality (left three plots).

To quantify how well the ranking based on the quality measures matches the ranking based on downstream performance, we compute the Spearman correlation  $\rho$  between these quantities. In Table 1 we can see that the eigenspace overlap score gets consistently higher correlation values with downstream performance than the other measures of compression quality. Note that  $\Delta_{\max}$  also attains relatively high correlation values, though the eigenspace overlap score still outperforms  $\Delta_{\max}$  by 0.06 to 0.24 on the tasks in Table 1. See Appendix E.5 for similar results on other tasks.

## 4.2 Downstream Performance of Uniform Quantization

We show that across tasks and compression rates uniform quantization consistently matches or outperforms the other compression methods, in terms of both the eigenspace overlap score and downstream performance. These empirical results validate our analysis from Section 3.2 showing that uniformly quantized embeddings in expectation attain high eigenspace overlap scores, and are thus likely to attain strong downstream performance. In Figure 3 we plot the average eigenspace overlap (left) and average question answering (SQuAD) performance (right) of compressed fastText embeddings for different compression methods and compression rates; we visualize the standard deviation over five random seeds with error bars. Our primary conclusion is that the simple uniform quantization method consistently performs similarly to or better than the other compression methods, both in terms of the eigenspace overlap score and downstream performance.<sup>7</sup> Given the connections between downstream performance and the eigenspace overlap score, the high eigenspace overlap scores attained by uniform quantization help explain its strong downstream performance. For results with the same trend on the GLUE and sentiment tasks, see Appendices E.1, E.4.<sup>8</sup>




Figure 3: **Eigenspace overlap and downstream performance of uniform quantization.** Uniform quantization can attain high values for the eigenspace overlap  $\mathcal{E}$ , and match the k-means and DCCL methods for fastText embeddings on the question answering (SQuAD) task.

<sup>7</sup>We apply uniform quantization to compress embeddings trained end-to-end for a translation task in Appendix E.2; we show it outperforms a tensorized factorization [16] proposed for the task-specific setting.

<sup>8</sup>We provide a memory-efficient implementation of the uniform quantization method in <https://github.com/HazyResearch/smallfry>.



Table 2: **The selection error rate of each measure of compression quality as a selection criterion.** Across all pairs of compressed embeddings from our experiments, we measure for each task the fraction of cases when a quality measure selects the worse performing embedding. We observe that the eigenspace overlap score  $\mathcal{E}$  achieves lower error rates than other compression quality measures.

Dataset	SQuAD		SST-1		MNLI	QQP
Embedding	GloVe	fastText	GloVe	fastText	BERT WordPiece	BERT WordPiece
PIP loss	0.32	0.37	0.32	0.40	0.31	0.32
$\Delta$	0.34	0.58	0.39	0.57	0.32	0.33
$\Delta_{\max}$	0.28	0.22	0.30	0.27	0.15	0.16
$1 - \mathcal{E}$	<b>0.17</b>	<b>0.11</b>	<b>0.19</b>	<b>0.20</b>	<b>0.10</b>	<b>0.10</b>

### 4.3 Compressed Embedding Selection with the Eigenspace Overlap Score

We now show that the eigenspace overlap score is a more accurate and robust selection criterion for compressed embeddings than the existing measures of compression quality. In our experiment, we first enumerate all the embeddings we compressed using different compression methods, compression rates, and five random seeds, and we evaluate each of these embeddings on the various downstream tasks; we use the same random seed for compression and for downstream training. We then consider for each task all pairs of compressed embeddings, and for each measure of compression quality report the *selection error rate*—the fraction of cases where the embedding with a higher compression quality score attains worse downstream performance. We show in Table 2 that across different tasks the eigenspace overlap score achieves lower selection error rates than the PIP loss and the spectral distance measures  $\Delta$  and  $\Delta_{\max}$ , with  $1.3\times$  to  $2\times$  lower selection error rates than the second best measure. To demonstrate the robustness of the eigenspace overlap score as a criterion, we measure the maximum difference in downstream performance, across all pairs of compressed embeddings discussed above, between the better performing embedding and the one selected by the eigenspace overlap score. We observe that this maximum performance difference is  $1.1\times$  to  $5.5\times$  smaller for the eigenspace overlap score than for the measure of compression quality with the second smallest maximum performance difference. See Appendix E.8 for more detailed results on the robustness of the eigenspace overlap score as a selection criterion.

## 5 Related Work

Compressing machine learning models is critical for training and inference in resource-constrained settings. To enable low-memory training, recent work investigates using low numerical precision [21, 9] and sparsity [35, 24]. To compress a model for low-memory inference, Han et al. [14] investigate pruning and quantization for deep neural networks.

Our work on understanding the generalization performance of compressed embeddings is also closely related to work on understanding the generalization performance of kernel approximation methods [38, 31]. In particular, training a linear model over compressed word embeddings can be viewed as training a model with a linear kernel using an approximation to the kernel matrix. Recently, there has been work on how different measures of kernel approximation error relate to the generalization performance of the model trained using the approximate kernels, with Avron et al. [3] and Zhang et al. [41] proposing the spectral measures of approximation error which we consider in this work.

## 6 Conclusion and Future Work

We proposed the eigenspace overlap score, a new way to measure the quality of a compressed embedding without requiring training for each downstream task of interest. We related this score to the generalization performance of linear and logistic regression models, used this score to better understand the strong empirical performance of uniformly quantized embeddings, and showed that this score is an accurate and robust selection criterion for compressed embeddings. Although this work focuses on word embeddings, for future work we hope to show that the ideas presented here extend to other domains—for example, to other types of embeddings (e.g., graph node embeddings [12]), and to compressing the activations of neural networks. We also believe that our work can help understand the performance of any model trained using compressed or perturbed features, and to understand why certain proposed methods for compressing neural networks succeed while others fail. We hope this work inspires improvements to compression methods in various domains.

## Acknowledgments

We thank Tony Ginart, Max Lam, Stephanie Wang, and Christopher Aberger for all their work on the early stages of this project. We further thank all the members of our research group for their helpful discussions and feedback throughout the course of this work.

We gratefully acknowledge the support of DARPA under Nos. FA87501720095 (D3M), FA86501827865 (SDH), and FA86501827882 (ASED); NIH under No. U54EB020405 (Mobilize), NSF under Nos. CCF1763315 (Beyond Sparsity), CCF1563078 (Volume to Velocity), and 1937301 (RTML); ONR under No. N000141712266 (Unifying Weak Supervision); the Moore Foundation, NXP, Xilinx, LETI-CEA, Intel, IBM, Microsoft, NEC, Toshiba, TSMC, ARM, Hitachi, BASF, Accenture, Ericsson, Qualcomm, Analog Devices, the Okawa Foundation, American Family Insurance, Google Cloud, Swiss Re, and members of the Stanford DAWN project: Teradata, Facebook, Google, Ant Financial, NEC, VMWare, and Infosys. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views, policies, or endorsements, either expressed or implied, of DARPA, NIH, ONR, or the U.S. Government.

## References

- [1] Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. Globally normalized transition-based neural networks. In *ACL*, 2016.
- [2] Martin Andrews. Compressing word embeddings. In *ICONIP*, 2016.
- [3] Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. Random Fourier features for kernel ridge regression: Approximation bounds and statistical guarantees. In *ICML*, 2017.
- [4] Nicola Bertoldi, Prashant Mathur, Nicholas Ruiz, and Marcello Federico. FBK’s machine translation and speech translation systems for the IWSLT 2014 evaluation campaign. In *IWSLT*, 2014.
- [5] Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. Reading Wikipedia to answer open-domain questions. In *ACL*, 2017.
- [6] Ting Chen, Martin Renqiang Min, and Yizhou Sun. Learning k-way d-dimensional discrete codes for compact embedding representations. In *ICML*, 2018.
- [7] Dan Shiebler, Chris Green, Luca Belli, Abhishek Tayal. Embeddings@Twitter, 2018. URL [https://blog.twitter.com/engineering/en\\_us/topics/insights/2018/embeddingsattwitter.html](https://blog.twitter.com/engineering/en_us/topics/insights/2018/embeddingsattwitter.html). [Online; published 13-Sept-2018; accessed 20-May-2019].
- [8] C. Davis and W. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- [9] Christopher De Sa, Megan Leszczynski, Jian Zhang, Alana Marzoev, Christopher R Aberger, Kunle Olukotun, and Christopher Ré. High-accuracy low-precision training. *arXiv preprint arXiv:1803.03383*, 2018.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [11] A. Gersho. Quantization. *IEEE Communications Society Magazine*, 15(5):16–16, Sep. 1977.
- [12] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *KDD*, 2016.
- [13] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *ICML*, 2015.
- [14] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding. In *ICLR*, 2016.

- [15] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- [16] Valentin Khrulkov, Oleksii Hrinchuk, Leyla Mirvakhabova, and Ivan V. Oseledets. Tensorized embedding layers for efficient model compression. *arXiv preprint arXiv:1901.10787*, 2019.
- [17] J. Kiefer. Sequential minimax search for a maximum. *Proceedings of the American Mathematical Society*, 4:502–506, 1953.
- [18] Yoon Kim. Convolutional neural networks for sentence classification. In *EMNLP*, 2014.
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [20] Omer Levy and Yoav Goldberg. Neural word embedding as implicit matrix factorization. In *NeurIPS*, 2014.
- [21] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Frederick Diamos, Erich Elsen, David García, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed precision training. In *ICLR*, 2018.
- [22] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [23] Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhresch, and Armand Joulin. Advances in pre-training distributed word representations. In *LREC*, 2018.
- [24] Hesham Mostafa and Xin Wang. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *ICML*, 2019.
- [25] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. fairseq: A fast, extensible toolkit for sequence modeling. In *NAACL-HLT: Demonstrations*, 2019.
- [26] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NeurIPS Autodiff Workshop*, 2017.
- [27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [28] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global vectors for word representation. In *EMNLP*, 2014.
- [29] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *NAACL-HLT*, 2018.
- [30] Tiberiu Popoviciu. Sur les équations algébriques ayant toutes leurs racines réelles. *Mathematica*, 9:129–145, 1935.
- [31] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NeurIPS*, 2007.
- [32] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *EMNLP*, 2016.
- [33] Raphael Shu and Hideki Nakayama. Compressing word embeddings via deep compositional code learning. In *ICLR*, 2018.
- [34] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*, 2013.
- [35] Nimit Sharad Sohoni, Christopher Richard Aberger, Megan Leszczynski, Jian Zhang, and Christopher Ré. Low-memory neural network training: A technical report. *arXiv preprint arXiv:1904.10631*, 2019.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

- [37] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*, 2019.
- [38] Christopher K. I. Williams and Matthias W. Seeger. Using the Nyström method to speed up kernel machines. In *NeurIPS*, 2000.
- [39] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.
- [40] Zi Yin and Yuanyuan Shen. On the dimensionality of word embedding. In *NeurIPS*, 2018.
- [41] Jian Zhang, Avner May, Tri Dao, and Christopher Ré. Low-precision random Fourier features for memory-constrained kernel approximation. In *AISTATS*, 2019.

## A Background

### A.1 Uniform Quantization

A  $b$ -bit uniform quantization of a real number  $x \in [-r, r]$  is computed as follows: First, the interval  $[-r, r]$  is divided into  $2^b - 1$  sub-intervals of equal size. Then,  $x$  is rounded to either the top or bottom of the sub-interval  $[\underline{x}, \bar{x}]$  containing  $x$ , where  $\underline{x} = r + j \frac{2r}{2^b-1}$  and  $\bar{x} = r + (j+1) \frac{2r}{2^b-1}$ , for  $j \in \{0, 1, \dots, 2^b - 2\}$ . Given this rounded value, one can simply store the  $b$ -bit integer  $j$  or  $j+1$  in place of the real-valued  $x$ , depending on whether  $x$  was rounded to  $\underline{x}$  or  $\bar{x}$  respectively. In this work, we will consider a deterministic rounding scheme which rounds  $x$  to the nearest value (denoted by  $Q_{b,r}(x)$ ), as well as an unbiased stochastic rounding scheme (denoted by  $\tilde{Q}_{b,r}(x)$ ). More details below). Note that our analysis will focus on the stochastic rounding scheme, while our experiments will focus on deterministic quantization; however, for completeness, in Appendix E.9, we show that stochastic quantization also performs quite well empirically.

We now define unbiased stochastic uniform quantization more formally. We will denote by  $\tilde{Q}_{b,r}(x)$  the  $b$ -bit unbiased stochastic uniform quantization of a real number  $x \in [-r, r]$ . More formally, if  $x \in [\underline{x}, \bar{x}]$  for  $\underline{x} = r + j \frac{2r}{2^b-1}$  and  $\bar{x} = r + (j+1) \frac{2r}{2^b-1}$ , for  $j \in \{0, 1, \dots, 2^b - 2\}$ ,  $\mathbb{P}[\tilde{Q}_{b,r}(x) = \underline{x}] = \frac{\bar{x}-x}{\bar{x}-\underline{x}}$  and  $\mathbb{P}[\tilde{Q}_{b,r}(x) = \bar{x}] = \frac{x-\underline{x}}{\bar{x}-\underline{x}}$ . Note that  $\mathbb{E}[\tilde{Q}_{b,r}(x)] = x$  and  $\text{VAR}[\tilde{Q}_{b,r}(x)] \leq \frac{r^2}{(2^b-1)^2} = \delta_b^2 r^2$  for  $\delta_b^2 := \frac{1}{(2^b-1)^2}$ . We bound the variance using the fact that a bounded random variable in an interval of length  $c$  has variance at most  $c^2/4$  by Popoviciu’s inequality on variances [30] (in our case,  $c = \frac{2r}{2^b-1}$ ).

Using the above definition of  $\tilde{Q}_{b,r}$ , we define the  $b$ -bit stochastic uniform quantization of a matrix  $X$ :

**Definition 2.** For a bounded embedding matrix  $X$  with  $X_{ij} \in [-r, r]$ , we define a  $b$ -bit stochastic uniform quantization of  $X$  to be a matrix  $\tilde{X}$  such that  $\tilde{X}_{ij} = \tilde{Q}_{b,r}(X_{ij})$ .

For details on how we use uniform quantization to compress word embeddings, please see Algorithm 1 and the associated discussion in Appendix D.3.

### A.2 Fixed Design Linear Regression

We derive here the close form expression for the risk of fixed design linear regression. In this setting we observe a set of labeled points  $\{(x_i, y_i)\}_{i=1}^n$  where the observe labels  $y_i = \bar{y}_i + \epsilon_i \in \mathbb{R}$  are perturbed versions of the true label  $\bar{y}$  with independent zero-mean noise  $\epsilon_i$  (with variance  $\sigma^2$ ). In other words,  $y = \bar{y} + \epsilon$  with  $\epsilon$  being a  $n$ -dimensional zero-mean random variable with covariance  $\sigma^2 I_n$ . Let  $X \in \mathbb{R}^{n \times d}$  be the feature matrix. The weight vector  $w^*$  of the optimal linear regressor  $f_{X,\epsilon}(x) = \langle x, w^* \rangle$  is computed by minimizing the least square loss:

$$w^* = \arg \min_{w \in d} \frac{1}{n} \|Xw - y\|^2.$$

From the normal equation, we know that  $w^* = (X^T X)^{-1} X^T y$ . The risk, or expected error, of the optimal linear regressor  $f_{X,\epsilon}$  trained on data matrix  $X$  and label vector  $y = \bar{y} + \epsilon$  is defined as

$$\begin{aligned}\mathcal{R}_{\bar{y}}(X) &:= \mathbb{E}_{\epsilon} \left[ \frac{1}{n} \sum_{i=1}^n (f_{X,\epsilon}(x_i) - \bar{y}_i)^2 \right] \\ &= \mathbb{E}_{\epsilon} \left[ \frac{1}{n} \|Xw^* - \bar{y}\|^2 \right].\end{aligned}$$

**Proposition 4.** *If the feature matrix  $X \in \mathbb{R}^{n \times d}$  is full-rank and has the SVD decomposition  $X = USV^T$ , then the risk of the optimal linear regressor, in the fixed design linear regression problem with noise variance  $\sigma^2$ , is*

$$\mathcal{R}_{\bar{y}}(X) = \frac{1}{n} (\|\bar{y}\|^2 - \|U^T \bar{y}\|^2 + d\sigma^2).$$

*Proof.* From the normal equation,

$$w^* = (X^T X)^{-1} X^T y = (VS^2V^T)^{-1} VSU^T y = VS^{-2}V^T VSU^T y = VS^{-2}SU^T y = VS^{-1}U^T y.$$

Substituting this expression into the definition of the risk, we obtain

$$\begin{aligned}\mathcal{R}_{\bar{y}}(X) &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|Xw - \bar{y}\|^2] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|USV^T VS^{-1}U^T y - \bar{y}\|^2] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|UU^T y - \bar{y}\|^2] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|UU^T(\bar{y} + \epsilon) - \bar{y}\|^2] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|(UU^T - I_n)(\bar{y} + \epsilon) + \epsilon\|^2] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|(I_n - UU^T)(\bar{y} + \epsilon) - \epsilon\|^2] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\|A(\bar{y} + \epsilon) - \epsilon\|^2] \quad (\text{letting } A := I_n - UU^T) \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [(\bar{y} + \epsilon)^T A^2 (\bar{y} + \epsilon) - (\bar{y} + \epsilon)^T A \epsilon - \epsilon^T A (\bar{y} + \epsilon) + \epsilon^T \epsilon] \\ &= \frac{1}{n} \mathbb{E}_{\epsilon} [\bar{y}^T A \bar{y} - \epsilon^T A \epsilon + \epsilon^T \epsilon] \quad (\text{using } A^2 = A, \text{ and } \mathbb{E}_{\epsilon} [\epsilon^T A \bar{y}] = \mathbb{E}_{\epsilon} [\bar{y}^T A \epsilon] = 0) \\ &= \frac{1}{n} \left( \bar{y}^T (I_n - UU^T) \bar{y} + \mathbb{E}_{\epsilon} [-\epsilon^T (I_n - UU^T) \epsilon + \epsilon^T \epsilon] \right) \\ &= \frac{1}{n} \left( \|\bar{y}\|^2 - \|U^T \bar{y}\|^2 + \mathbb{E}_{\epsilon} [\epsilon^T UU^T \epsilon] \right) \\ &= \frac{1}{n} \left( \|\bar{y}\|^2 - \|U^T \bar{y}\|^2 + d\sigma^2 \right),\end{aligned}$$

where the last step follows from

$$\mathbb{E}_{\epsilon} [\epsilon^T UU^T \epsilon] = \mathbb{E}_{\epsilon} [\text{tr}(UU^T \epsilon \epsilon^T)] = \text{tr}(UU^T \mathbb{E}_{\epsilon} [\epsilon \epsilon^T]) = \text{tr}(UU^T \sigma^2 I_n) = \sigma^2 \|U\|_F^2 = d\sigma^2.$$

□

## B The Eigenspace Overlap Score: Theory and Extensions

### B.1 Proof of Theorem 1: Average Case Analysis for Fixed Design Linear Regression

We present the proof of Theorem 1, relating the generalization performance and eigenspace overlap score in the context of fixed design linear regression. The true label  $\bar{y}$  is assumed to be randomly

distributed in the span of  $U$ , of the form  $\bar{y} = Uz$  for some zero-mean  $d$ -dimensional random variable  $z$ . While in Section 3 we assume for simplicity that  $z$  has identity covariance, here we consider the more general setting of  $z$  having covariance matrix  $\Sigma$ . Note that because  $\frac{1}{n}\mathbb{E}[\|\bar{y}\|^2] = \frac{1}{n}\mathbb{E}_z[\|Uz\|^2] = \frac{1}{n}\mathbb{E}_z[\|z\|^2] = \frac{1}{n}\text{tr}(\Sigma)$ , to maintain a constant signal to noise ratio it makes sense for the variance  $\sigma^2$  of the noise we add to each entry of  $y$  to scale as  $\sigma^2 = O(\frac{1}{n}\text{tr}(\Sigma))$ . Thus, we introduce a scalar  $c \in \mathbb{R}$  such that  $\sigma^2 = \frac{c^2}{n}\text{tr}(\Sigma)$ ; this is the more general form of  $\sigma^2 = c^2 \frac{d}{n}$  from Section 3. We now prove the more general version of Theorem 1.

**Theorem 1 (Generalized).** *Let  $X = USV^T \in \mathbb{R}^{n \times d}$  be the singular value decomposition of a full-rank embedding matrix  $X$ , and let  $\tilde{X} \in \mathbb{R}^{n \times k}$  be another full-rank embedding matrix. Let  $\bar{y} = Uz \in \mathbb{R}^n$  denote a random label vector in  $\text{span}(U)$ , where  $z \in \mathbb{R}^d$  has mean zero and covariance matrix  $\Sigma$ . Let  $\lambda_{\min}(\Sigma)$  be the smallest eigenvalue of  $\Sigma$ . Letting  $\sigma^2 = \frac{c^2}{n}\text{tr}(\Sigma)$  denote the variance of the label noise, it follows that*

$$\mathbb{E}_{\bar{y}} \left[ \mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X) \right] \leq \frac{\text{tr} \Sigma - d\lambda_{\min}(\Sigma)\mathcal{E}(X, \tilde{X})}{n} - c^2 \cdot \frac{\text{tr}(\Sigma)(d-k)}{n^2}.$$

If the random vector  $z$  has identity covariance matrix, then

$$\mathbb{E}_{\bar{y}} \left[ \mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X) \right] = \frac{d}{n} \cdot \left( 1 - \mathcal{E}(X, \tilde{X}) \right) - c^2 \cdot \frac{d(d-k)}{n^2}.$$

*Proof.* Since  $\bar{y} = Uz$ , we have

$$\begin{aligned} \mathbb{E}_{\bar{y}} [\|U^T \bar{y}\|^2] &= \mathbb{E}_z [\|U^T Uz\|^2] = \mathbb{E}_z [\|z\|^2] = \mathbb{E}_z [\text{tr}(z^T z)] \\ &= \mathbb{E}_z [\text{tr}(zz^T)] = \text{tr}(\mathbb{E}_z [zz^T]) = \text{tr}(\Sigma). \end{aligned}$$

Similarly,

$$\begin{aligned} \mathbb{E}_{\bar{y}} [\|\tilde{U}^T \bar{y}\|^2] &= \mathbb{E}_z [\|\tilde{U}^T Uz\|^2] = \mathbb{E}_z [\text{tr}(z^T U^T \tilde{U} \tilde{U}^T Uz)] = \mathbb{E}_z [\text{tr}(U^T \tilde{U} \tilde{U}^T U z z^T)] \\ &= \text{tr}(U^T \tilde{U} \tilde{U}^T U \mathbb{E}_z [z z^T]) = \text{tr}(U^T \tilde{U} \tilde{U}^T U \Sigma) = \text{tr}(\Sigma^{1/2} U^T \tilde{U} \tilde{U}^T U \Sigma^{1/2}) \\ &= \|\tilde{U}^T U \Sigma^{1/2}\|_F^2, \end{aligned}$$

where  $\Sigma^{1/2}$  is the positive semidefinite (PSD) matrix such that  $(\Sigma^{1/2})^2 = \Sigma$ . From Proposition 4, the risks are  $\mathcal{R}_{\bar{y}}(X) = \frac{1}{n}(\|\bar{y}\|^2 - \|U^T \bar{y}\|^2 + d\sigma^2)$  and  $\mathcal{R}_{\bar{y}}(\tilde{X}) = \frac{1}{n}(\|\bar{y}\|^2 - \|\tilde{U}^T \bar{y}\|^2 + k\sigma^2)$ . We thus obtain:

$$\begin{aligned} \mathbb{E}_{\bar{y}} [\mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X)] &= \frac{1}{n} \left( \mathbb{E}_{\bar{y}} [\|U^T \bar{y}\|^2] - \mathbb{E}_{\bar{y}} [\|\tilde{U}^T \bar{y}\|^2] - (d-k)\sigma^2 \right) \\ &= \frac{1}{n} \left( \text{tr}(\Sigma) - \|\tilde{U}^T U \Sigma^{1/2}\|_F^2 \right) - \frac{d-k}{n} \sigma^2. \end{aligned} \quad (2)$$

We can lower bound  $\|\tilde{U}^T U \Sigma^{1/2}\|_F^2$  in terms of the smallest eigenvalue of  $\Sigma$  and the eigenspace overlap score of  $X$  and  $\tilde{X}$ . Specifically, we now show that  $\|\tilde{U}^T U \Sigma^{1/2}\|_F^2 \geq d\lambda_{\min}(\Sigma)\mathcal{E}(X, \tilde{X})$ , where  $\lambda_{\min}(\Sigma)$  is the smallest eigenvalue of  $\Sigma$ . We will use the fact that  $\Sigma - \lambda_{\min}(\Sigma)I_d$  is PSD, and so  $(\Sigma - \lambda_{\min}(\Sigma)I_d)^{1/2}$  exists. We now prove the above inequality:

$$\begin{aligned} \|\tilde{U}^T U \Sigma^{1/2}\|_F^2 - d\lambda_{\min}(\Sigma)\mathcal{E}(X, \tilde{X}) &= \|\tilde{U}^T U \Sigma^{1/2}\|_F^2 - \lambda_{\min}(\Sigma)\|\tilde{U}^T U\|_F^2 \\ &= \text{tr}(\Sigma^{1/2} U^T \tilde{U} \tilde{U}^T U \Sigma^{1/2}) - \lambda_{\min}(\Sigma) \text{tr}(U^T \tilde{U} \tilde{U}^T U) \\ &= \text{tr}(U^T \tilde{U} \tilde{U}^T U \Sigma) - \text{tr}(U^T \tilde{U} \tilde{U}^T U \lambda_{\min}(\Sigma)I_d) \\ &= \text{tr}\left(U^T \tilde{U} \tilde{U}^T U (\Sigma - \lambda_{\min}(\Sigma)I_d)\right) \\ &= \text{tr}\left((\Sigma - \lambda_{\min}(\Sigma)I_d)^{1/2} U^T \tilde{U} \tilde{U}^T U (\Sigma - \lambda_{\min}(\Sigma)I_d)^{1/2}\right) \\ &= \|\tilde{U}^T U (\Sigma - \lambda_{\min}(\Sigma)I_d)^{1/2}\|_F^2 \\ &\geq 0. \end{aligned}$$

Thus, we have shown that  $\mathbb{E}_{\bar{y}} [\|\tilde{U}^T \bar{y}\|^2] = \|\tilde{U}^T U \Sigma^{1/2}\|_F^2 \geq d \lambda_{\min}(\Sigma) \mathcal{E}(X, \tilde{X})$ . Substituting this lower bound into Equation (2) yields

$$\mathbb{E}_{\bar{y}} [\mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X)] \leq \frac{1}{n} \left( \text{tr}(\Sigma) - d \lambda_{\min}(\Sigma) \mathcal{E}(X, \tilde{X}) \right) - \frac{d-k}{n} \sigma^2.$$

In the case where  $\Sigma = I_d$ , we obtain  $\text{tr}(\Sigma) = d$  and  $\|\tilde{U}^T U \Sigma^{1/2}\|_F^2 = \|\tilde{U}^T U\|_F^2 = d \mathcal{E}(X, \tilde{X})$ . Thus from Equation (2), we obtain

$$\mathbb{E}_{\bar{y}} [\mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X)] = \frac{d}{n} (1 - \mathcal{E}(X, \tilde{X})) - \frac{d-k}{n} \sigma^2.$$

Substituting  $\sigma^2 = \frac{c^2}{n} \text{tr}(\Sigma)$  in the above expressions, and noting that  $\text{tr}(\Sigma) = d$  when  $\Sigma = I_d$ , completes the proof.  $\square$

## B.2 Average Case Analysis for Lipschitz-Continuous Loss Function

We now consider the fixed design setting with a Lipschitz-continuous loss function, and discuss how the average risk of training on  $\tilde{X}$  can be bounded in terms of the average risk of training on  $X$  and the eigenspace overlap score  $\mathcal{E}(X, \tilde{X})$ .

Let  $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  be a non-negative loss function which is  $L$ -Lipschitz in both its first and second arguments,  $X \in \mathbb{R}^{n \times d}$  be a fixed data matrix with SVD  $X = U S V^T$ ,  $x_i \in \mathbb{R}^d$  be the  $i^{\text{th}}$  row of  $X$ , and  $y \in \mathbb{R}^n$  be a label vector. We assume that  $\arg \min_v \ell(v', v) = v$  for all  $v \in \mathbb{R}$ . We will consider a linear model  $f(x) = x^T w$  parameterized by some weight vector  $w$ , such that the loss function for each data point  $x_i$  under this model is  $\ell(x_i^T w, y_i)$ .

Similar to the fixed design linear regression setting, we assume that  $y$  is generated from the true label  $\bar{y} \in \mathbb{R}^n$  by adding zero-mean independent noise:  $y = \bar{y} + \epsilon$  where  $\epsilon = [\epsilon_1, \dots, \epsilon_n]^T \in \mathbb{R}^n$  and the  $\epsilon_i$  are independent with zero mean and variance  $\sigma^2$ . We also assume that the true label  $\bar{y}$  is randomly distributed in the span of  $U$ , of the form  $\bar{y} = U z$  for some zero-mean  $d$ -dimensional random variable  $z$ , with covariance matrix  $\Sigma$ . Lastly, we let  $\tilde{X} = \tilde{U} \tilde{S} \tilde{V}^T \in \mathbb{R}^{n \times k}$  be any matrix, which for our purposes will represent a compressed version of  $X$ .

We now define the optimal weight vectors  $w^*$  and  $\tilde{w}^*$  trained on  $X$  and  $\tilde{X}$  respectively, along with their corresponding vectors of predictions  $u$  and  $\tilde{u}$ :

$$\begin{aligned} w^* &= \arg \min_w \sum_{i=1}^n \ell(x_i^T w, y_i), \quad u := X w^* \\ \tilde{w}^* &= \arg \min_{\tilde{w}} \sum_{i=1}^n \ell(\tilde{x}_i^T \tilde{w}, y_i), \quad \tilde{u} := \tilde{X} \tilde{w}^*. \end{aligned} \quad (3)$$

Note that  $u$  and  $\tilde{u}$  depend on  $\epsilon$  and  $\bar{y}$ , which are random.

The risks, or expected errors (expectation taken over  $\epsilon$ , for a fixed  $\bar{y}$ ) for the models trained with  $X$  and  $\tilde{X}$  respectively, are defined as

$$\mathcal{R}_{\bar{y}}(X) := \mathbb{E}_{\epsilon} \left[ \frac{1}{n} \sum_{i=1}^n \ell(x_i^T w^*, \bar{y}_i) \right]. \quad (4)$$

$$\mathcal{R}_{\bar{y}}(\tilde{X}) := \mathbb{E}_{\epsilon} \left[ \frac{1}{n} \sum_{i=1}^n \ell(\tilde{x}_i^T \tilde{w}^*, \bar{y}_i) \right]. \quad (5)$$

We are now ready to present our Theorem for the case of average-case generalization performance with  $L$ -Lipschitz continuous loss functions.

**Theorem 2 (Generalized).** *Let  $X \in \mathbb{R}^{n \times d}$ ,  $\tilde{X} \in \mathbb{R}^{n \times k}$ ,  $\bar{y} \in \mathbb{R}^n$ ,  $z \in \mathbb{R}^d$ ,  $\Sigma \in \mathbb{R}^{d \times d}$ ,  $\lambda_{\min} \in \mathbb{R}$ , and  $c \in \mathbb{R}$  be defined as in Theorem 1 (Generalized). Let  $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  be a convex non-negative loss function which is  $L$ -Lipschitz continuous in both arguments and satisfies  $\arg \min_v \ell(v', v) = v \forall v \in \mathbb{R}$ . It follows that*

$$\mathbb{E}_{\bar{y}} [\mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X)] \leq \frac{L}{\sqrt{n}} \left( \sqrt{\text{tr}(\Sigma) - d \lambda_{\min}(\Sigma) \mathcal{E}(X, \tilde{X})} + 2c \sqrt{\text{tr}(\Sigma)} \right).$$

If the random vector  $z$  has identity covariance matrix, then

$$\mathbb{E}_{\bar{y}} \left[ \mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X) \right] \leq \frac{L\sqrt{d}}{\sqrt{n}} \left( \sqrt{1 - \mathcal{E}(X, \tilde{X})} + 2c \right).$$

*Proof.* Let  $f(v, y) = \frac{1}{n} \sum_{i=1}^n \ell(v_i, y_i)$  be the average loss on the training set, given predictions  $v$ , and let  $f(v, \bar{y}) = \frac{1}{n} \sum_{i=1}^n \ell(v_i, \bar{y}_i)$  be the average test loss. Note that  $f$  is  $L/\sqrt{n}$ -Lipschitz in its first argument, because  $\|\nabla_u f(v, y)\|^2 \leq \sum_{i=1}^n (L/n)^2 = L^2/n$ . Similarly,  $f$  is  $L/\sqrt{n}$ -Lipschitz in its second argument. The training loss on  $X$  is then  $f(u, y)$ , and the risk is  $\mathbb{E}_\epsilon [f(u, \bar{y})]$ . Similarly, the training loss on  $\tilde{X}$  is  $f(\tilde{u}, y)$  and the risk is  $\mathbb{E}_\epsilon [f(\tilde{u}, \bar{y})]$ .

We can bound the difference in the average risk (average over  $\bar{y}$ ) when training on  $X$  and  $\tilde{X}$  in terms of the eigenspace overlap score. We do this in three steps: First, we lower bound  $\mathcal{R}_{\bar{y}}(X)$ . Second, we upper bound  $\mathcal{R}_{\bar{y}}(\tilde{X})$ . Third, we used the bounds from the first two steps to upper bound the expectation over  $\bar{y}$  of the difference between  $\mathcal{R}_{\bar{y}}(X)$  and  $\mathcal{R}_{\bar{y}}(\tilde{X})$ . We now go through these steps one at a time:

- **Step 1:** We show that  $\mathcal{R}_{\bar{y}}(X) \geq f(\bar{y}, \bar{y})$ .

$$\mathcal{R}_{\bar{y}}(X) = \mathbb{E}_\epsilon [f(u, \bar{y})] \geq \mathbb{E}_\epsilon [f(\bar{y}, \bar{y})] = f(\bar{y}, \bar{y}).$$

Here, we used the fact that  $\bar{y} = \arg \min_v f(v, \bar{y})$  (which follows from our assumption on the loss function  $\ell$ ).

- **Step 2:** We show that for all  $\tilde{w} \in \mathbb{R}^k$ ,  $\mathcal{R}_{\bar{y}}(\tilde{X}) \leq f(\tilde{X}\tilde{w}, \bar{y}) + 2L\sigma$ .

$$\begin{aligned} \mathcal{R}_{\bar{y}}(\tilde{X}) &= \mathbb{E}_\epsilon [f(\tilde{u}, \bar{y})] \\ &\leq \mathbb{E}_\epsilon \left[ f(\tilde{u}, y) + \frac{L}{\sqrt{n}} \|y - \bar{y}\| \right] \quad (f \text{ is } L/\sqrt{n}\text{-Lipschitz}) \\ &= \mathbb{E}_\epsilon [f(\tilde{X}\tilde{w}^*, y)] + \mathbb{E}_\epsilon \left[ \frac{L}{\sqrt{n}} \|\epsilon\| \right] \\ &\leq \mathbb{E}_\epsilon [f(\tilde{X}\tilde{w}, y)] + \mathbb{E}_\epsilon \left[ \frac{L}{\sqrt{n}} \|\epsilon\| \right] \quad (\text{by Equation (3)}) \\ &\leq \mathbb{E}_\epsilon \left[ f(\tilde{X}\tilde{w}, \bar{y}) + \frac{L}{\sqrt{n}} \|\bar{y} - y\| \right] + \mathbb{E}_\epsilon \left[ \frac{L}{\sqrt{n}} \|\epsilon\| \right] \quad (f \text{ is } L/\sqrt{n}\text{-Lipschitz}) \\ &= f(\tilde{X}\tilde{w}, \bar{y}) + \frac{2L}{\sqrt{n}} \mathbb{E}_\epsilon [\|\epsilon\|] \\ &\leq f(\tilde{X}\tilde{w}, \bar{y}) + 2L\sigma \quad (\text{by } \mathbb{E}_\epsilon [\|\epsilon\|] \leq \sqrt{\mathbb{E}_\epsilon [\|\epsilon\|^2]} = \sqrt{n\sigma^2}). \end{aligned}$$

- **Step 3:** We bound the expected difference, over the randomness in the label vector  $\bar{y}$ , between  $\mathcal{R}_{\bar{y}}(\tilde{X})$  and  $\mathcal{R}_{\bar{y}}(X)$ , leveraging the results from steps 1 and 2 above.

$$\begin{aligned} \mathbb{E}_{\bar{y}} \left[ \mathcal{R}_{\bar{y}}(\tilde{X}) - \mathcal{R}_{\bar{y}}(X) \right] &\leq \mathbb{E}_{\bar{y}} \left[ f(\tilde{X}\tilde{w}, \bar{y}) + 2L\sigma - f(\bar{y}, \bar{y}) \right] \quad (\text{by steps 1 and 2}) \\ &\leq \mathbb{E}_{\bar{y}} \left[ \frac{L}{\sqrt{n}} \|\tilde{X}\tilde{w} - \bar{y}\| + 2L\sigma \right] \quad (f \text{ is } L/\sqrt{n}\text{-Lipschitz}). \end{aligned}$$

To get the tightest bound, we can minimize  $\|\tilde{X}\tilde{w} - \bar{y}\|$  over  $\tilde{w} \in \mathbb{R}^k$ . But this is exactly the least squares problem, with solution  $\tilde{w} = (\tilde{X}^T \tilde{X})^{-1} \tilde{X}^T \bar{y} = \tilde{V} \tilde{S}^{-1} \tilde{U}^T \bar{y}$ , and minimum value  $\sqrt{\|\bar{y}\|^2 - \|\tilde{U}^T \bar{y}\|^2}$  (by proof of Proposition 4). We can substitute this bound into the above



inequalities and continue:

$$\begin{aligned}
\mathbb{E}_{\tilde{y}} [\mathcal{R}_{\tilde{y}}(\tilde{X}) - \mathcal{R}_{\tilde{y}}(X)] &\leq \frac{L}{\sqrt{n}} \mathbb{E}_{\tilde{y}} \left[ \sqrt{\|\tilde{y}\|^2 - \|\tilde{U}^T \tilde{y}\|^2} \right] + 2L\sigma \\
&\leq \frac{L}{\sqrt{n}} \sqrt{\mathbb{E}_{\tilde{y}} [\|\tilde{y}\|^2 - \|\tilde{U}^T \tilde{y}\|^2]} + 2L\sigma \quad (\text{by Jensen's inequality}) \\
&\leq \frac{L}{\sqrt{n}} \sqrt{\text{tr}(\Sigma) - d\lambda_{\min}(\Sigma)\mathcal{E}(X, \tilde{X})} + 2L\sigma,
\end{aligned}$$

where this last step follows from  $\mathbb{E}_{\tilde{y}} [\|\tilde{y}\|^2] = \mathbb{E}_z [z^T U^T U z] = \mathbb{E}_z [z^T z] = \text{tr}(\Sigma)$ , and  $\mathbb{E}_{\tilde{y}} [\|\tilde{U}^T \tilde{y}\|^2] \geq d\lambda_{\min}(\Sigma)\mathcal{E}(X, \tilde{X})$ , which we show in the proof of Theorem 1 (Generalized).

Using the assumption  $\sigma^2 = \frac{c^2}{n} \text{tr}(\Sigma)$ , and thus  $\sigma = \frac{c\sqrt{\text{tr}(\Sigma)}}{\sqrt{n}}$ , completes the proof.  $\square$

Note that in the case of logistic regression where we observe the noisy logits,<sup>9</sup> the loss is 1-Lipschitz in the first argument. If we assume that the weight vector  $w$  has bounded norm (say, because of L2 regularization), and that the data matrix  $X$  is bounded, then the loss function is also Lipschitz in the second argument. We can think of  $z_i$  as being the optimal logits such that  $P(y_i = 1) = \sigma(z_i)$  (one can think of these  $z_i = x_i^T w$  as the parameters of the generative model which generated the data). Just like in the linear regression case, we see that the overlap  $\mathcal{E}(X, \tilde{X}) = \frac{1}{d} \|\tilde{U}^T U\|_F^2$  gives an upper bound on the maximum possible expected difference in the loss functions when training on  $X$  vs.  $\tilde{X}$ .

### B.3 Robustness of the Eigenspace Overlap Score to Perturbations

For a measure of compression quality to correlate strongly with downstream performance, a necessary condition is for it to be robust to embedding perturbations which are unlikely to significantly affect generalization performance. Here, we give an example of an embedding perturbation which has minimal effect on the eigenspace overlap score and on average-case generalization performance, while having a much larger impact on the other measures of compression quality. We consider the following simple perturbation: if  $X = \sum_{i=1}^d \sigma_i U_i V_i^T$  is the singular value decomposition of  $X$ , we consider setting its largest singular value to 0, resulting in the perturbed matrix  $\tilde{X} := \sum_{i=2}^d \sigma_i U_i V_i^T$ . Assuming a label vector  $y = Uz$ ,  $\tilde{X}$  would have generalization error of  $\|y\|^2 - \|\tilde{U}^T y\|^2 = z_1^2$ . If we assume that  $z_1^2 \ll \sum_{i=2}^d z_i^2$  (as would be expected in our average-case analysis), then  $\tilde{X}$  would perform similarly to  $X$ .

In Table 3, we show the impact of the above perturbation on the various measures of compression quality we have discussed. At a high-level, we observe that this perturbation can have a dramatic effect of the previously proposed measures, while having minimal effect on the eigenspace overlap score. For example, the eigenspace overlap score after this perturbation is equal to  $\mathcal{E}(X, \tilde{X}) = \frac{d-1}{d}$ , relative to the maximum possible overlap of 1. In contrast, this perturbation results in a  $\Delta_1$  value very close to 1 if  $\lambda \ll \sigma_1$  (note that 1 is the maximum possible value for  $\Delta_1$ , and that Zhang et al. [41] show generalization bounds scale with  $\frac{1}{1-\Delta_1}$ ). This makes sense, because  $\Delta_1$  can be used to attain a worst-case generalization bound for the perturbed embeddings, and there exist cases where setting the largest singular value to 0 can significantly harm the generalization performance of the embeddings (e.g., if  $z_1^2 \approx \|z\|^2$ ). Thus, while the  $\Delta_1$  measure is important for understanding the worst-case performance of the compressed embeddings, it is generally an overly pessimistic measure. The eigenspace overlap score, on the other hand, is generally unable to provide worst-case guarantees, but aligns nicely with the expected performance of the compressed embeddings in the average-case setting.

<sup>9</sup>For logistic regression, we can write  $\ell(z', z) := -(\sigma(z) \log(\sigma(z')) + (1 - \sigma(z)) \log(1 - \sigma(z')))$ . Here  $z$  and  $z'$  both represent logits. We can recover the standard logistic loss by letting  $z \in \{-\infty, \infty\}$ . Or, you can think of  $\sigma(z)$  as  $p(y = 1 | x)$ , the parameter of the Bernoulli generating the label for a datapoint  $x$ .

Table 3: **Effect of perturbation on measures of compression quality.** In this table, we consider the effect of perturbing an embedding matrix  $X$  by setting its largest singular value to 0 on the various measures of compression quality discussed above. As we can see, setting the largest singular value of  $X$  to 0 can have a disproportionately large effect on the relative reconstruction error  $\left(\frac{\|X - \tilde{X}\|_F}{\|X\|_F}\right)$ , relative PIP loss  $\left(\frac{\|XX^T - \tilde{X}\tilde{X}^T\|_F}{\|XX^T\|_F}\right)$ ,  $\Delta_1$ ,  $\Delta$ , and  $\Delta_{\max}$  measures (values can approach 1), while having a modest effect on the eigenspace overlap score (value of  $1/d$  always).

Compression quality measure	Measure after perturbation
Rel. reconstruction error	$\sigma_1 / \sqrt{\sum_{i=1}^d \sigma_i^2}$
Rel. PIP loss	$\sigma_1^2 / \sqrt{\sum_{i=1}^d \sigma_i^4}$
$\Delta_1$	$\sigma_1^2 / (\sigma_1^2 + \lambda)$
$\Delta_2$	0
$\Delta$	$\sigma_1^2 / (\sigma_1^2 + \lambda)$
$\Delta_{\max}$	$(\sigma_1^2 + \lambda) / \lambda$
$1 - \mathcal{E}(X, \tilde{X})$	$\frac{1}{d}$

#### B.4 Relating the Eigenspace Overlap Score to Embedding Reconstruction Error

We now define a variant of embedding reconstruction error which we show is closely related to the eigenspace overlap score. As we mention in Section 2.2, the definition of embedding reconstruction error  $\|X - \tilde{X}\|_F$  is only applicable when  $X \in \mathbb{R}^{n \times d}$  and  $\tilde{X} \in \mathbb{R}^{n \times k}$  have the same dimensions ( $d = k$ ). To get around this limitation, we define the *projected embedding reconstruction error* as  $\min_{P \in \mathbb{R}^{k \times d}} \|\tilde{X}P - X\|_F^2$ . It is easy to show that the matrix  $P$  minimizing the above expression is  $P^* := (\tilde{X}^T \tilde{X})^{-1} \tilde{X}^T X$ . Letting  $X = USV^T$  and  $\tilde{X} = \tilde{U}\tilde{S}\tilde{V}^T$  be the singular value decompositions of  $X$  and  $\tilde{X}$ , we can simplify the expression for the projected embedding reconstruction error as follows:

$$\begin{aligned}
\min_{P \in \mathbb{R}^{k \times d}} \|\tilde{X}P - X\|_F^2 &= \|\tilde{X}(\tilde{X}^T \tilde{X})^{-1} \tilde{X}^T X - X\|_F^2 \\
&= \|\tilde{U}\tilde{S}\tilde{V}^T(\tilde{V}\tilde{S}^{-2}\tilde{V}^T)\tilde{V}S\tilde{U}^T X - X\|_F^2 \\
&= \|\tilde{U}\tilde{U}^T X - X\|_F^2 \\
&= \text{tr} \left( (\tilde{U}\tilde{U}^T X - X)^T (\tilde{U}\tilde{U}^T X - X) \right) \\
&= \|X\|_F^2 - \|\tilde{U}^T X\|_F^2 \\
&= \|X\|_F^2 - \|\tilde{U}^T USV^T\|_F^2 \\
&= \|X\|_F^2 - \|\tilde{U}^T US\|_F^2
\end{aligned}$$

Thus, the projected embedding reconstruction error is equal to a term  $(\|X\|_F^2)$  which is constant in  $\tilde{X}$ , minus a term  $\|\tilde{U}^T US\|_F^2 = \sum_{i=1}^d \sigma_i^2 \|\tilde{U}^T U_i\|_2^2$ . Note that this second term is simply a version of the eigenspace overlap score  $\frac{1}{\max(d,k)} \|\tilde{U}^T U\|_F^2 = \frac{1}{\max(d,k)} \sum_{i=1}^d \|\tilde{U}^T U_i\|_2^2$  which weights the projections of the different singular vectors  $U_i$  of  $X$  onto  $\tilde{U}$  according to the singular values of  $X$ . In Section B.1 we show that in the case where the random label vector  $\bar{y} = Uz$  where  $z$  is a zero mean random variable in  $\mathbb{R}^d$  with covariance  $\Sigma$ , the expected error depends on a term  $\|\tilde{U}^T U \Sigma^{1/2}\|_F^2$ . Thus, the projected embedding reconstruction error is directly related to the expected error when  $z$  is sampled with covariance matrix  $\Sigma = S^2$ .

In Table 4 we show that the projected embedding reconstruction error, like the eigenspace overlap score, attains high Spearman correlation with downstream performance.

Table 4: **Spearman correlation between projected embedding reconstruction error and downstream performance.** In this table, we show that the Spearman correlation  $\rho$  between the projected embedding reconstruction error and downstream performance is relatively similar to the Spearman correlation between the eigenspace overlap score and downstream performance. We show results on the SQuAD question answering task, and the SST-1 sentiment analysis task, for both GloVe and fastText embeddings. In each table entry, we present the correlation absolute values as “GloVe  $|\rho|$  | fastText  $|\rho|$ .”

	SQuAD		SST-1	
Projected embed. reconst. error	0.82	0.86	0.75	0.64
$1 - \mathcal{E}$	0.81	0.91	0.75	0.73

## C The Eigenspace Overlap Score of Uniformly Quantized Embeddings

This Appendix focuses on the eigenspace overlap score of uniformly quantized embeddings. In Appendix C.1 we prove our result on the expected eigenspace overlap score of uniformly quantized embeddings (Theorem 3). In Appendix C.2 we validate that the empirical scaling of the eigenspace overlap score with respect to the vocabulary size, embedding dimension, compression rate, and smallest singular value of the embedding matrix, matches the scaling predicted by the theory. Lastly, in Appendix C.3, we demonstrate that choosing the clipping value for uniform quantization is crucial for attaining a high eigenspace overlap score, and that choosing the clipping threshold with lowest reconstruction error is very similar to choosing the clipping threshold with highest eigenspace overlap score. Additionally, we demonstrate that the optimal clipping thresholds for deterministic and stochastic quantization are very similar, and that deterministic quantization attains slightly higher eigenspace overlap scores than stochastic quantization.

### C.1 Theorem 3 Proof

We now prove Theorem 3, which bounds the expected eigenspace overlap scores for uniformly quantized embeddings. The core of our proof is an application of the Davis-Kahan  $\sin(\Theta)$  theorem [8]. We now review this classic theorem, and then prove our result.

**Theorem 5.** (*Davis-Kahan  $\sin(\Theta)$  Theorem (adapted)*) Let  $K = U_0 S_0 U_0^T + U_1 S_1 U_1^T$  be the eigendecomposition of  $K$  such that  $U_0 \in \mathbb{R}^{n \times d}$  are the first  $d$  eigenvectors of  $K = U S U^T$ ,  $S_0$  the first  $d$  eigenvalues,  $U_1, S_1$  the rest. Similarly, let  $\tilde{K} = V_0 R_0 V_0^T + V_1 R_1 V_1^T$  be the equivalent eigendecomposition for  $\tilde{K} = K + H$ . If the eigenvalues of  $S_0$  are contained in the interval  $(a_0, a_1)$ , and the eigenvalues of  $R_1$  are excluded from the interval  $(a_0 - \delta, a_1 + \delta)$  for some  $\delta > 0$ , then

$$\|V_1^T U_0\| \leq \frac{\|V_1^T H U_0\|}{\delta} \quad (6)$$

for any unitarily invariant norm  $\|\cdot\|$ .

To prove Theorem 3, we will apply the Davis-Kahan  $\sin(\Theta)$  theorem to the setting where  $K$  is the Gram matrix of an uncompressed matrix  $X$ , and  $\tilde{K}$  is the gram matrix of a  $b$ -bit stochastic uniform quantization  $\tilde{X}$  of  $X$  (See Definition 2). We now present and prove Theorem 3.

**Theorem 3.** Let  $X \in \mathbb{R}^{n \times d}$  be a bounded embedding matrix with  $X_{ij} \in [-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}]$  and smallest singular value  $\sigma_{\min} = a\sqrt{n/d}$ , for  $a \in (0, 1]$ .<sup>10</sup> Let  $\tilde{X}$  be a  $b$ -bit stochastic uniform quantization of  $X$ . Then for  $n \geq \max(33, d)$ , we can lower bound the expected eigenspace overlap score of  $\tilde{X}$ , over the randomness of the stochastic quantization, as follows:

$$\mathbb{E} [1 - \mathcal{E}(X, \tilde{X})] \leq \frac{20}{(2^b - 1)^2 a^4}.$$

*Proof.* We will denote the Gram matrices of  $X$  and  $\tilde{X}$  by  $K = X X^T = U S U^T$  and  $\tilde{K} = \tilde{X} \tilde{X}^T = (X + C)(X + C)^T = V R V^T$ . Here,  $C$  is a stochastic matrix satisfying  $\mathbb{E}[C_{ij}] = 0$

<sup>10</sup>The maximum possible value of  $\sigma_{\min}$  is  $\sqrt{n/d}$ , which occurs when  $\|X\|_F^2 = n$  and  $\sigma_{\min} = \sigma_{\max}$ .

and  $\text{VAR}[C_{ij}] \leq \delta_b^2/d \quad \forall i, j$ , for  $\delta_b^2 := \frac{1}{(2^b-1)^2}$  (see Appendix A.1). In our application of the Davis-Kahan  $\sin(\Theta)$  theorem, we will use  $a_0 = \sigma_{\min}(K)$ ,  $a_1 = \infty$ ,  $\delta = \sigma_{\min}(K)$ . Note also the  $H = \tilde{K} - K = (X + C)(X + C)^T - XX^T = XC^T + CX^T + CC^T$ . We will let  $a \in [0, 1]$  be the scalar such that  $\sigma_{\min}(X) = a\sqrt{\frac{n}{d}}$  (equivalently,  $\sigma_{\min}(K) = a^2\frac{n}{d}$ ).

Using the Davis-Kahan  $\sin(\Theta)$  theorem, along with Lemma 6 (below), we can show the following:

$$\begin{aligned}
\|V_1^T U_0\|_F &\leq \frac{\|V_1^T H U_0\|_F}{\sigma_{\min}(K)} \\
&= \frac{\|V_1^T (XC^T + CX^T + CC^T) U_0\|_F}{\sigma_{\min}(K)} \\
&\leq \frac{\|V_1^T\|_2 \|XC^T + CX^T + CC^T\|_F \|U_0\|_2}{\sigma_{\min}(K)} \quad (\text{using } \|AB\|_F \leq \|A\|_2 \|B\|_F \text{ twice.}) \\
&\leq \frac{\|XC^T + CX^T + CC^T\|_F}{\sigma_{\min}(K)} \quad (\text{using } \|V_1^T\|_2 = \|U_0\|_2 = 1.) \\
\Rightarrow \frac{1}{d} \|V_1^T U_0\|_F^2 &\leq \frac{\|XC^T + CX^T + CC^T\|_F^2}{d \cdot \sigma_{\min}(K)^2} \\
\iff 1 - \frac{1}{d} \|V_0^T U_0\|_F^2 &\leq \frac{\|XC^T + CX^T + CC^T\|_F^2}{d \cdot \sigma_{\min}(K)^2} \quad (\text{using } \|V_0^T U_0\|_F^2 + \|V_1^T U_0\|_F^2 = \|U_0\|_F^2 = d) \\
\iff 1 - \mathcal{E}(X, \tilde{X}) &\leq \frac{\|XC^T + CX^T + CC^T\|_F^2}{d \cdot \sigma_{\min}(K)^2}
\end{aligned}$$

$$\begin{aligned}
\Rightarrow \mathbb{E} [1 - \mathcal{E}(X, \tilde{X})] &\leq \mathbb{E} \left[ \frac{\|XC^T + CX^T + CC^T\|_F^2}{d \cdot \sigma_{\min}(K)^2} \right] \\
&= \frac{\mathbb{E} [\|XC^T + CX^T + CC^T\|_F^2]}{d \cdot \sigma_{\min}(K)^2} \\
&\leq \frac{\frac{20n^2\delta_b^2}{d}}{d \cdot \sigma_{\min}(K)^2} \quad (\text{by Lemma 6.}) \\
&= \frac{20n^2\delta_b^2}{d^2 a^4 (n^2/d^2)} \\
&= \frac{20\delta_b^2}{a^4}
\end{aligned}$$

□

We now present and prove Lemma 6.

**Lemma 6.** *Let  $X \in \mathbb{R}^{n \times d}$  be a bounded embedding matrix with  $X_{ij} \in [-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}]$ . Let  $\tilde{X} = X + C$  be a  $b$ -bit stochastic uniform quantization of  $X$ . Then for  $n \geq \max(33, d)$ , it follows that*

$$\mathbb{E} [\|XC^T + CX^T + CC^T\|_F^2] \leq \frac{20n^2\delta_b^2}{d}. \quad (7)$$

*Proof.* We will let  $H := XC^T + CX^T + CC^T$ . To bound  $\mathbb{E} [\|H\|_F^2] = \sum_{i,j=1}^n \mathbb{E} [H_{ij}^2]$ , we will consider two cases:  $H_{ij}$  for  $i \neq j$  and  $H_{ij}$  for  $i = j$ . We will let  $x_i, c_i \in \mathbb{R}^d$  denote the  $i^{\text{th}}$  rows of  $X$  and  $C$  respectively.

1. **Case 1:**  $i \neq j$

$$\begin{aligned}
\mathbb{E}[H_{ij}^2] &= \mathbb{E}[(x_i^T c_j + c_i^T x_j + c_i^T c_j)^2] \\
&= \mathbb{E}[(x_i^T c_j)^2 + (c_i^T x_j)^2 + (c_i^T c_j)^2] \\
&= \mathbb{E}\left[\left(\sum_{k=1}^d x_{ik} c_{jk}\right)^2\right] + \mathbb{E}\left[\left(\sum_{k=1}^d c_{ik} x_{jk}\right)^2\right] + \mathbb{E}\left[\left(\sum_{k=1}^d c_{ik} c_{jk}\right)^2\right] \\
&= \mathbb{E}\left[\sum_{k=1}^d x_{ik}^2 c_{jk}^2\right] + \mathbb{E}\left[\sum_{k=1}^d c_{ik}^2 x_{jk}^2\right] + \mathbb{E}\left[\sum_{k=1}^d c_{ik}^2 c_{jk}^2\right] \\
&= \sum_{k=1}^d x_{ik}^2 \mathbb{E}[c_{jk}^2] + \sum_{k=1}^d \mathbb{E}[c_{ik}^2] x_{jk}^2 + \sum_{k=1}^d \mathbb{E}[c_{ik}^2] \mathbb{E}[c_{jk}^2] \\
&\leq \frac{\delta_b^2}{d} \cdot \sum_{k=1}^d x_{ik}^2 + \frac{\delta_b^2}{d} \cdot \sum_{k=1}^d x_{jk}^2 + \sum_{k=1}^d \left(\frac{\delta_b^2}{d}\right)^2 \\
&\leq \frac{\delta_b^2}{d} \cdot \|x_i\|^2 + \frac{\delta_b^2}{d} \cdot \|x_j\|^2 + \sum_{k=1}^d \left(\frac{\delta_b^2}{d}\right)^2 \\
&\leq \frac{2\delta_b^2 + \delta_b^4}{d} \quad (\text{using } \|x_i\|^2 \leq 1) \\
&\leq \frac{3\delta_b^2}{d} \quad (\text{using } \delta_b \leq 1).
\end{aligned}$$

2. **Case 2:**  $i = j$

$$\begin{aligned}
\mathbb{E}[H_{ii}^2] &= \mathbb{E}[(x_i^T c_i + c_i^T x_i + c_i^T c_i)^2] \\
&= \mathbb{E}\left[\left(2 \sum_{k=1}^d x_{ik} c_{ik} + \sum_{l=1}^d c_{il}^2\right)^2\right] \\
&= \mathbb{E}\left[4 \left(\sum_{k=1}^d x_{ik} c_{ik}\right)^2 + 4 \left(\sum_{k=1}^d x_{ik} c_{ik}\right) \cdot \left(\sum_{l=1}^d c_{il}^2\right) + \left(\sum_{l=1}^d c_{il}^2\right)^2\right] \\
&= \mathbb{E}\left[4 \sum_{k=1}^d x_{ik}^2 c_{ik}^2 + 4 \sum_{k,l=1}^d x_{ik} c_{ik} c_{il}^2 + \sum_{k,l=1}^d c_{il}^2 c_{ik}^2\right] \\
&= 4 \sum_{k=1}^d x_{ik}^2 \mathbb{E}[c_{ik}^2] + 4 \sum_{k=1}^d x_{ik} \mathbb{E}[c_{ik}^3] + \sum_{k,l=1}^d \mathbb{E}[c_{il}^2 c_{ik}^2] \\
&\leq 4 \cdot \frac{\delta_b^2}{d} \cdot \sum_{k=1}^d x_{ik}^2 + 4 \sum_{k=1}^d \frac{1}{\sqrt{d}} \left(\frac{2}{\sqrt{d}(2^b-1)}\right)^3 + \sum_{k,l=1}^d \left(\frac{2}{\sqrt{d}(2^b-1)}\right)^4 \\
&= 4 \cdot \frac{\delta_b^2}{d} \cdot \|x_i\|^2 + 4d \cdot \frac{8}{d^2(2^b-1)^3} + d^2 \cdot \frac{16}{d^2(2^b-1)^4} \\
&\leq \frac{4\delta_b^2}{d} + \frac{32}{d(2^b-1)^3} + \frac{16}{(2^b-1)^4} \\
&= \frac{4\delta_b^2 + 32\delta_b^3}{d} + 16\delta_b^4 \\
&\leq \frac{36\delta_b^2}{d} + 16\delta_b^4 \quad (\text{using } \delta_b \leq 1).
\end{aligned}$$

Now we can combine the above results:

$$\begin{aligned}
\sum_{i,j=1}^n \mathbb{E}[H_{ij}^2] &\leq \sum_{i \neq j} \left( \frac{3\delta_b^2}{d} \right) + \sum_{i=1}^n \left( \frac{36\delta_b^2}{d} + 16\delta_b^4 \right) \\
&= n(n-1) \left( \frac{3\delta_b^2}{d} \right) + n \left( \frac{36\delta_b^2}{d} + 16\delta_b^4 \right) \\
&= \frac{3n^2\delta_b^2 - 3n\delta_b^2 + 36n\delta_b^2}{d} + 16n\delta_b^4 \\
&= \frac{3n^2\delta_b^2 + 33n\delta_b^2}{d} + 16n\delta_b^4 \\
&\leq \frac{4n^2\delta_b^2}{d} + 16n\delta_b^4 \quad (\text{assuming } n \geq 33.) \\
&\leq \frac{4n^2\delta_b^2}{d} + \frac{16n^2\delta_b^2}{d} \quad (\text{assuming } n \geq d.) \\
&= \frac{20n^2\delta_b^2}{d}
\end{aligned}$$

□

## C.2 Empirical Validation of Theorem 3 Scaling

We now validate Theorem 3 empirically by showing the impact of the precision ( $b$ ), the scalar ( $a$ ), the vocabulary size ( $n$ ), and the embedding dimension ( $d$ ) on the eigenspace overlap score  $\mathcal{E}(X, \tilde{X})$  of uniformly quantized embeddings matrices. As predicted by the theory, we will show in Figure 4 that  $1 - \mathcal{E}(X, \tilde{X})$  drops as  $b$  and  $a$  are increased, and is relatively unaffected by changes in  $n$  and  $d$ .

We now describe our experimental protocol for studying the impact of each of these parameters on the eigenspace overlap score:

- **Precision ( $b$ ), Figure 4(a):** We randomly generate a  $10^4 \times 10$  matrix, with entries drawn uniformly from  $[-\frac{1}{\sqrt{10}}, \frac{1}{\sqrt{10}}]$ . We uniformly quantize this matrix with precisions  $b \in \{1, 2, 4, 8, 16\}$ , and compute the eigenspace overlap score between the quantized matrix and the original matrix. As one can see,  $1 - \mathcal{E}(X, \tilde{X})$  drops rapidly as the precision is increased.
- **Scalar ( $a$ ), Figure 4(b):** We randomly generate a  $10^4 \times 10$  matrix, with entries drawn uniformly from  $[-\frac{1}{\sqrt{10}}, \frac{1}{\sqrt{10}}]$ . We then multiply this matrix on the right by diagonal matrices with diagonal entries spaced logarithmically between 1 and  $\{1, 0.1, .01, .001, .0001\}$ , thus generating matrices with increasingly small values of the scalar  $a$ . We uniformly quantize each of these matrices with precisions  $b \in \{1, 2, 4\}$ , and compute the eigenspace overlap score between the quantized matrices and the original matrices. As one can see,  $1 - \mathcal{E}(X, \tilde{X})$  drops as the scalar  $a$  increases.
- **Vocabulary size ( $n$ ), Figure 4(c):** We randomly generate  $n \times 10$  matrices for  $n \in \{10^2, 3 \times 10^2, 10^3, 3 \times 10^3, 10^4, 3 \times 10^4, 10^5\}$ , with entries drawn uniformly from  $[-\frac{1}{\sqrt{10}}, \frac{1}{\sqrt{10}}]$ . We uniformly quantize these matrices with precisions  $b \in \{1, 2, 4\}$ , and compute the corresponding eigenspace overlap scores. As one can see, the vocabulary size  $n$  has minimal impact on the eigenspace overlap score.
- **Embedding dimension ( $d$ ), Figure 4(d):** We randomly generate  $10^4 \times d$  matrices for  $d \in \{10, 30, 100, 300, 1000\}$  with entries drawn uniformly from  $[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}]$ . We uniformly quantize these matrices with precisions  $b \in \{1, 2, 4\}$ , and compute the corresponding eigenspace overlap scores. As one can see, the embedding dimension  $d$  has minimal impact on the eigenspace overlap score.

An important thing to mention about Theorem 3 is that this bound can be vacuous when the embedding matrix has a quickly decaying spectrum, and thus a small value of  $a$ . This is a consequence of the proof of the Davis-Kahan  $\sin(\Theta)$  theorem, which uses the smallest eigenvalue of  $XX^T$  to lower bound a matrix multiplication; this inequality is relatively tight when the spectrum of  $XX^T$  decays slowly, but is quite loose if it doesn't.




Figure 4: **Empirical Validation of Theorem 3.** We measure the eigenspace overlap score  $\mathcal{E}$  of uniformly quantized embeddings with the uncompressed embedding for various precisions  $b$  and scalar  $a$ , vocabulary sizes  $n$ , and dimensions  $d$ . We observe that  $1 - \mathcal{E}$  decays as the precision  $b$  and scalar  $a$  grow, and that  $1 - \mathcal{E}$  is largely unaffected by the vocabulary size  $n$  and embedding dimension  $d$ .




Figure 5: **The impact of clipping and deterministic vs. stochastic quantization on the eigenspace overlap score.** (a) We plot the eigenspace overlap score as a function of the clipping threshold, for precisions  $b \in \{1, 2, 4\}$  and for both stochastic and deterministic quantization. We observe that choosing the value of  $r$  appropriately is crucial for attaining high eigenspace overlap scores, and that deterministic quantization generally gives slightly higher eigenspace overlap scores than stochastic quantization. (b) For each precision  $b$ , we plot the clipping thresholds which give the highest eigenspace overlap scores and embedding reconstruction errors, for both stochastic and deterministic quantization. We observe that for both types of quantization, the optimal clipping threshold chosen according to embedding reconstruction error is very similar to the optimal clipping threshold chosen according to the eigenspace overlap score.

### C.3 Impact of Clipping and Deterministic vs. Stochastic Quantization on the Eigenspace Overlap Score

As shown in Algorithm 1 (described in Section D.3), clipping is the first step in the uniform quantization method we use for compressing word embeddings. Here, we show that clipping is important because it can significantly improve the eigenspace overlap scores of the compressed embeddings, compared to uniform quantization without clipping. Specifically, we compute the eigenspace overlap score of  $Q_{b,r}(\text{clip}_r(X))$  (and  $\tilde{Q}_{b,r}(\text{clip}_r(X))$ ) with  $X$ , for a range of clipping values  $r \in [0, \max(|X|)]$ , using the publicly available 300-dimensional pre-trained GloVe embeddings as  $X$  (see Appendix D.2 for embedding details). Recall that  $Q_{b,r}$  and  $\tilde{Q}_{b,r}$  are the deterministic and stochastic  $b$ -bit uniform quantization functions for the interval  $[-r, r]$ , respectively (defined in Section A.1). In Figure 5(a), we plot the eigenspace overlap scores attained by both quantization methods as a function of the clipping value  $r$ , for precisions  $b \in \{1, 2, 4\}$ . We observe that choosing the value of  $r$  appropriately is crucial for attaining high eigenspace overlap scores. We also observe that deterministic quantization typically attains slightly higher eigenspace overlap scores than stochastic quantization. This result helps explain our empirical observation in Appendix E.9 that deterministic quantization often attains slightly better downstream performance than stochastic quantization.

In Algorithm 1, we choose the clipping threshold  $r^*$  which minimizes the embedding reconstruction error of the clipped and quantized embeddings. In Figure 5(b), we show that choosing the clipping threshold based on the embedding reconstruction error gives very similar results to choosing the

clipping threshold based on the eigenspace overlap score, for both deterministic and stochastic quantization. This helps explain the strong downstream performance of the embeddings compressed using Algorithm 1.

## D Experiment Details

We now discuss in detail the protocols we used for all our experiments. In Appendix D.1, we describe the model architectures and datasets we use for each downstream task, including the train/development/test splits for each dataset. We then discuss in Appendix D.2 the details of the pre-trained word embeddings we compress, and in Appendix D.3 the details of the different compression methods we use. In Appendix D.4 we discuss the training details for each of the downstream tasks, including the hyperparameter grids we use to tune our models.

### D.1 Task Details

**Question Answering** For the question answering task, we use the DrQA model [5] trained and evaluated on the Stanford Question and Answering Dataset (SQuAD) [32]. For this task, given a paragraph and a corresponding question in natural language, the model must predict the start and end position, within the paragraph, of the answer to the question. We use the default train and development set splits for the SQuAD dataset, and report all results on the development set, as the test set is not publicly available. The DrQA model consists of a three-layer bidirectional LSTM model with 128-dimensional hidden units on top of a pretrained word embedding. We train the DrQA model on the SQuAD-v1.1 training set, and report the F1 score on the SQuAD-v1.1 development set. We use the implementation of the DrQA model from the Facebook Research DrQA repository.<sup>11</sup>

**Sentiment Analysis** For the sentiment analysis tasks, we use the convolutional neural network (CNN) architecture proposed by Kim [18], and evaluate performance on the datasets used in that work (see Section 3 of that paper for dataset details). We use the data released as part of the Harvard NLP group’s sentiment analysis repository.<sup>12</sup> For the datasets which are pre-split into train/development/test (SST-1, SST-2), we use these dataset splits. For the datasets which are pre-split into train/test (TREC), we take a random 10% of the training set as a development set. For the datasets which have no pre-specified splits (MR, Subj, CR, MPQA), we take a random 10% of the data as a test set, and a random 10% of the remaining data as a development set; the rest of the data is used as the training set. We tune hyperparameters (learning rate) on the development sets, and report results on the test sets. The CNN architecture we use for this task has one convolutional layer with multiple filters, followed by a ReLU non-linearity and a max-pooling layer. The convolutional layer uses filter windows of size 3, 4, and 5, each with 100 feature maps. As we use PyTorch [26] for all our experiments, we reimplemented this model architecture in PyTorch, using the original Theano implementation as a template.<sup>13</sup>

**GLUE Tasks** The General Language Understanding Evaluation (GLUE) benchmark [37] is a collection of nine natural language understanding tasks. We summarize these tasks in Table 5, along with the evaluation metric used for each task. We use the default train and development set splits for each of these tasks. We tune hyperparameters (learning rate) on the development sets, and also report results on the development sets, as the test sets are not publicly available. For each task, we use the standard approach of adding a linear layer on top of the pre-trained BERT model, and then fine-tuning the model using the data for that task. To evaluate the performance of compressed embeddings on these tasks, we compress the WordPiece [39] embeddings in the pre-trained case-sensitive BERT<sub>BASE</sub> model, and then fine-tune all the non-embedding model parameters, keeping the embeddings frozen during training. We use a third-party implementation of the BERT model, and of the fine-tuning procedure.<sup>14</sup> We run experiments on all the GLUE tasks except WNLI. We skip the WNLI dataset because this is a dataset on which it is very difficult to outperform the trivial model which always

<sup>11</sup><https://github.com/facebookresearch/DrQA>.

<sup>12</sup>[https://github.com/harvardnlp/sent\\_conv\\_torch/tree/master/data](https://github.com/harvardnlp/sent_conv_torch/tree/master/data).

<sup>13</sup>[https://github.com/yoonkim/CNN\\_sentence](https://github.com/yoonkim/CNN_sentence).

<sup>14</sup>PyTorch implementation of the pre-trained BERT model: <https://github.com/huggingface/pytorch-pretrained-BERT>. We use the examples/run\_classifier.py file provided in this repo for fine-tuning.



Table 5: **The GLUE datasets, along with the evaluation metric used for each dataset.** For the MRPC and QQP datasets, the average of the F1 score and accuracy on the development set is used. For the STS-B dataset, the average of the Pearson and Spearman correlations on the development set is used. For the MNLI dataset, the average of the accuracies on the matched and mismatched development sets is used.

Datasets	Evaluation Metrics
The Corpus of Linguistic Acceptability (CoLA)	Matthew’s Correlation
The Stanford Sentiment Treebank (SST-2)	Accuracy
Microsoft Research Paraphrase Corpus (MRPC)	F1 / Accuracy
Semantic Textual Similarity Benchmark (STS-B)	Pearson-Spearman Correlation
Quora Question Pairs (QQP)	F1 / Accuracy
Multi-Genre Natural Language Inference (MNLI)	Accuracy (matched/mismatched)
Question Natural Language Inference (QNLI)	Accuracy
Recognizing Textual Entailment (RTE)	Accuracy
Winograd Natural Language Inference (WNLI)	Accuracy

Table 6: The optimal learning rates  $\eta$  and dictionary sizes  $k$  for DCCL.

Embedding	GloVe			fastText			BERT WordPiece		
Compression rate	8×	16×	32×	8×	16×	32×	8×	16×	32×
$k$	8	4	4	8	4	8	128	64	32
$\eta$	0.0003	0.0003	0.0003	0.0001	0.0001	0.0001	0.0003	0.0003	0.0003

outputs the majority class. This trivial model attains 65.1% accuracy, and only two of the contributors to the GLUE leaderboard<sup>15</sup> have outperformed this model, as of this writing.

## D.2 Word Embedding Details

For the GloVe embeddings, we use publicly available embeddings pre-trained on the Wikipedia 2014 and Gigaword 5 corpora.<sup>16</sup> These are available for dimensions  $d \in \{50, 100, 200, 300\}$ ; we use the 300-dimensional embeddings for all our experiments, except for our GloVe dimensionality reduction experiments, where we use the lower-dimensional embeddings. For the fastText embeddings, we use the publicly available 300-dimensional embeddings trained on the Wikipedia 2017 corpus, the UMBC webbase corpus, and the statmt.org news dataset.<sup>17</sup> For the WordPiece embeddings [39], we use the embeddings which are part of the pre-trained case-sensitive BERT<sub>BASE</sub> model, available through the Hugging Face BERT repository.<sup>18</sup>

## D.3 Compression Method Details

**Deep Compositional Code Learning (DCCL)** We give an overview of the DCCL method [33] in Section 2.1. The important hyperparameters for this method include the learning rate  $\eta$  of the Adam optimizer [19], the number of dictionaries  $m$ , the size  $k$  of each dictionary, the temperature parameter  $\tau$  for Gumbel sampling, and the mini-batch size. To select the learning rate  $\eta$  and the dictionary size  $k$  for each compression rate, we perform a grid search using the Cartesian product of  $\eta \in \{0.00001, 0.00003, 0.0001, 0.0003, 0.001\}$  and  $k \in \{2, 4, 8, 16\}$  for each uncompressed embedding type (GloVe, fastText, BERT WordPiece embeddings) and compression rate. Note that given a compression rate and a dictionary size  $k$ , this uniquely determines the number of dictionaries  $m$  to use. We select the combination of learning rate and dictionary size which minimizes the reconstruction error of the compressed embeddings. When compressing BERT WordPiece embedding, we extended the dictionary size grid to  $k \in \{2, 4, 8, 16, 32, 64, 128, 256\}$  to avoid the

<sup>15</sup><https://gluebenchmark.com/leaderboard/>

<sup>16</sup><http://nlp.stanford.edu/data/glove.6B.zip>.

<sup>17</sup><https://s3-us-west-1.amazonaws.com/fasttext-vectors/wiki-news-300d-1M.vec.zip>.

<sup>18</sup><https://github.com/huggingface/pytorch-pretrained-BERT>.

---

**Algorithm 1** Uniform quantization for word embeddings

---

- 1: **Input:** Embedding  $X \in \mathbb{R}^{n \times d}$ ; quantization func.  $Q_{b,r}$ ; clipping func.  $\text{clip}_r: \mathbb{R} \rightarrow [-r, r]$ .
  - 2: **Output:** Quantized embedding  $\tilde{X} \in \mathbb{R}^{n \times d}$ .
  - 3:  $r^* := \arg \min_{r \in [0, \max(|X|)]} \|Q_{b,r}(\text{clip}_r(X)) - X\|_F$ .
  - 4: **Return:**  $Q_{b,r^*}(\text{clip}_{r^*}(X))$ .
- 

optimal dictionary size touching the boundary of the grid. We provide the optimal learning rates and dictionary sizes in Table 6 for reproducibility. For the temperature parameter  $\tau$ , we follow Shu and Nakayama [33] and consistently use  $\tau = 1.0$ . For all our experiments we use a mini-batch size of 64, which is the default value in the DCCL repository.<sup>19</sup>

**K-means** The k-means clustering method can be used to compress embeddings as follows: First, the one-dimensional k-means clustering algorithm is run on all the scalar entries in the full-precision embedding matrix  $X$ . Then, each entry in  $X$  is replaced by the centroid to which it is closest. If  $2^b$  centroids are used during the clustering step, then for each entry of the compressed embedding matrix, only the integer  $j \in \{0, 1, \dots, 2^b - 1\}$  of the corresponding centroid needs to be stored; this requires  $b$  bits per entry. In our experiments, we use the Scikit Learn [27] implementation of k-means. We use the default configuration from Scikit Learn, which runs for a maximum of 300 iterations and can early stop if the relative decrease of the loss function is smaller than  $10^{-4}$ .

**Dimensionality Reduction** The two dimensionality reduction methods we consider are (1) using pre-trained lower-dimensional embeddings, and (2) principal component analysis (PCA). For the GloVe embeddings, we use the publicly available lower-dimensional embeddings described in Appendix D.2. These embeddings are available for dimensions  $d \in \{50, 100, 200, 300\}$ , where we consider the 300-dimensional embeddings to be the “uncompressed” embeddings. For our experiments with fastText and BERT WordPiece embeddings, we use PCA to reduce the dimension of the embeddings, as these embeddings are not publicly available in lower dimensions. When we compress the 300-dimensional fastText and GloVe embeddings with dimensionality reduction, we use compression rates in  $\{1\times, 1.5\times, 3\times, 6\times\}$ . For the 768-dimensional BERT WordPiece embeddings, we use compression rates in  $\{1\times, 2\times, 4\times, 8\times\}$ .

We now give details on how we implement the PCA dimensionality reduction method. For an embedding  $X \in \mathbb{R}^{n \times d}$  with vocabulary size  $n$  and dimension  $d$ , let  $X = USV^T$  be the SVD of  $X$  with  $U = [U_1, \dots, U_d]$ ,  $S = \text{diag}([s_1, \dots, s_d])$ , and  $V = [V_1, \dots, V_d]$ . If we let  $U_{(k)} := [U_1, \dots, U_k]$ ,  $S_{(k)} := \text{diag}([s_1, \dots, s_k])$ , and  $V_{(k)} := [V_1, \dots, V_k]$  then we use  $\tilde{X} := U_{(k)}S_{(k)}$  as the  $k$ -dimensional compressed embedding. Note that for the GLUE tasks, we instead use  $\tilde{X} := U_{(k)}S_{(k)}V_{(k)}^T$  to ensure that these compressed embeddings are compatible with the parameters of the pre-trained BERT model; because the dimension  $k$  of these compressed embeddings is small compared to the vocabulary size  $n$ , storing  $V_{(k)}^T$  requires a relatively small amount of additional memory.

**Uniform Quantization** In Algorithm 1 we show how we use uniform quantization to compress word embeddings. The input to the algorithm is an embedding matrix  $X \in \mathbb{R}^{n \times d}$ , where  $n$  is the size of the vocabulary, and  $d$  is the dimension of the embeddings. We define the function  $\text{clip}_r(x) = \max(\min(x, r), -r)$  for any non-negative  $r$ ; when matrices are passed in as inputs to this function, it clips the entries in an element-wise fashion. Given an input embedding and a desired numbers of bits to use per entry of the compressed embedding matrix, the uniform quantization method operates in two steps:

- **Step 1:** We find the value of  $r \in [0, \max(|X|)]$  which minimizes the reconstruction error of the quantized embeddings after  $X$  is clipped to  $[-r, r]$ . More formally, we let  $r^* := \arg \min_{r \in [0, \max(|X|)]} \|Q_{b,r}(\text{clip}_r(X)) - X\|_F$ , and use this value  $r^*$  to clip  $X$ . In our experiments, we find  $r^*$  to within a specified tolerance  $\epsilon = 0.01$  using the golden-section search algorithm [17]. To avoid stochasticity impacting the search process for the clipping threshold, we always use deterministic rounding in the search for  $r^*$ , regardless of whether we use stochastic rounding or deterministic nearest rounding in the final quantization after clipping the extremal values.
- **Step 2:** We quantize the clipped embeddings to  $b$  bits per entry with  $Q_{b,r}$ .

---

<sup>19</sup><https://github.com/zomux/neuralcompressor>.

Table 7: Training hyperparameter for DrQA on the SQuAD dataset.

Hyperparameter	Value
Optimizer	Adamax
Decay rates for 1st moment $\beta_1$	0.9
Decay rates for 2nd moment $\beta_2$	0.999
Adamax $\epsilon$	$10^{-8}$
Learning rate	$2 \times 10^{-3}$
Batchsize	32
Training epochs	40
Dropout	0.4

Table 8: Training hyperparameter shared across sentiment analysis datasets.

Hyperparameter	Value
Optimizer	Adam
Decay rates for 1st moment $\beta_1$	0.9
Decay rates for 2nd moment $\beta_2$	0.999
Adam $\epsilon$	$10^{-8}$
Batchsize	32
Training epochs	100
Dropout	0.5

In all of our main experiments on the downstream performance (question answering, sentiment analysis, GLUE tasks) of compressed word embeddings, we use the deterministic quantization function  $Q_{b,r}$  introduced in Appendix A.1 for both steps of this algorithm. However, in Appendix E.9 we use the stochastic quantization function  $\tilde{Q}_{b,r}$  for the second step of this compression algorithm, and show that it performs similarly to deterministic quantization on downstream tasks.

#### D.4 Training Details

We now discuss the training details for the different tasks we consider, focusing on how we tune the hyperparameters.

**Question Answering** We use the default hyperparameters from the Facebook Research DrQA implementation for all our question answering experiments, as these are tuned for the SQuAD dataset.<sup>20</sup> We summarize these hyperparameters in Table 7.

**Sentiment Analysis** We tune the learning rate for each of the sentiment analysis datasets using the grid  $\{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1.0\}$ . For this tuning process, we use the uncompressed embedding for each dataset and embedding type (GloVe, fastText), and pick the learning rate which attains highest average accuracy on the development set across five random seeds. This learning rate is then used to train the models that use the uncompressed embeddings, as well as the embeddings compressed using uniform quantization, k-means, and DCCL. Note that we tune the learning rate individually for each embedding compressed using dimensionality reduction (for both GloVe and fastText). We do this to ensure that the lower dimensionality of these compressed embeddings does not result in the learning rate being improperly tuned. We list the hyperparameters shared across datasets in Table 8 and list the optimal learning rate for each dataset and embedding type in Table 9.

**GLUE Tasks** We tune the learning rate for each of the GLUE tasks using the grid  $\{10^{-5}, 2 \times 10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 10^{-4}\}$ . When tuning the learning rate, we use the uncompressed WordPiece embeddings, and we fine-tune the entire model, without freezing the embedding parameters. For each task we pick the learning rate which gives the best average performance (according to the metrics in Table 5) on the development set, across five random seeds. The optimal learning rates are listed in Table 10 for all the GLUE tasks we run. We use the default values (from both the Google Research

<sup>20</sup><https://github.com/facebookresearch/DrQA>.

Table 9: The optimal learning rate  $\eta$  for different sentiment analysis datasets.

Datasets	MR	SST-1	SST-2	Subj	TREC	CR	MPQA
GloVe uncompressed	0.001	0.001	0.001	0.001	0.001	0.001	0.001
GloVe dim. red. $1\times$	0.001	0.001	0.001	0.001	0.001	0.001	0.001
GloVe dim. red. $1.5\times$	0.001	0.001	0.001	0.001	0.001	0.001	0.001
GloVe dim. red. $3\times$	0.0001	0.001	0.001	0.001	0.001	0.001	0.001
GloVe dim. red. $6\times$	0.001	0.001	0.001	0.001	0.001	0.001	0.001
fastText uncompressed	0.001	0.001	0.001	0.001	0.001	0.001	0.001
fastText dim. red. $1\times$	0.0001	0.001	0.001	0.001	0.001	0.001	0.001
fastText dim. red. $1.5\times$	0.0001	0.001	0.001	0.001	0.001	0.001	0.001
fastText dim. red. $3\times$	0.0001	0.001	0.001	0.001	0.001	0.001	0.001
fastText dim. red. $6\times$	0.001	0.0001	0.001	0.001	0.001	0.001	0.001

Table 10: The optimal learning rate  $\eta$  for different GLUE tasks

Tasks	MNLI	QQP	QNLI	SST-2	CoLA	STS-B	MRPC	RTE
$\eta$	$3 \times 10^{-5}$	$3 \times 10^{-5}$	$3 \times 10^{-5}$	$3 \times 10^{-5}$	$10^{-5}$	$2 \times 10^{-5}$	$2 \times 10^{-5}$	$2 \times 10^{-5}$

TensorFlow BERT repository<sup>21</sup> and the Hugging Face PyTorch BERT repository<sup>22</sup>) for the other hyperparameters. Specifically, we fine-tune the model for 3 epochs using the Adam optimizer with a mini-batch size of 32, and a weight decay strength of 0.01 (weight decay is not applied to the layer norm layers or to the bias parameters). We use a linear learning rate warm-up for the first 10% of training (learning rate grows linearly from  $0\times$  to  $1\times$  the specified learning rate), and then a linear learning rate decay for the remaining 90% of training (learning rate decays linearly from  $1\times$  to  $0\times$  the specified learning rate).

## D.5 Infrastructure Details

We run our experiments using AWS p2.xlarge instances, which have NVIDIA Tesla K80 GPUs. We use Python 3.6 for our experiments. For compatibility with the DrQA repository (which had not been ported to PyTorch 1.0 when we began our experiments), we use PyTorch 0.3.1 for the question answering and sentiment analysis tasks. For the GLUE tasks we use PyTorch 1.0.

## E Extended Empirical Results

We now provide a more complete version of the empirical results included in the main body of the paper, as well as a number of additional experiments validating claims related to our work. More specifically:

- In Appendix E.1 we present extended results comparing the downstream performance of the different compression methods across a range of compression rates for the GloVe, fastText, and BERT WordPiece embeddings, on question answering, sentiment analysis, and GLUE tasks. We show that uniform quantization can consistently match or outperform the other compression methods across these settings.
- In Appendix E.2 we present experiments comparing the performance of the different compression methods when applied to compressing task-specific embeddings which have been trained end-to-end for a translation task. Though our main focus in this paper is compressing task-agnostic embeddings (e.g., GloVe, fastText), we show that uniform quantization can effectively compete with a recently proposed tensorized factorization [16] of the embedding matrix designed for the task-specific setting.
- In Appendix E.3 we study whether, under a fixed memory budget, it is better to use low-dimensional high-precision embeddings, or high-dimensional low-precision embeddings.

<sup>21</sup>[https://github.com/google-research/bert/blob/master/run\\_classifier.py](https://github.com/google-research/bert/blob/master/run_classifier.py).

<sup>22</sup>[https://github.com/huggingface/pytorch-pretrained-BERT/blob/master/examples/run\\_classifier.py](https://github.com/huggingface/pytorch-pretrained-BERT/blob/master/examples/run_classifier.py).

We show that under a wide range of memory budgets, one can attain large improvements in downstream performance on the SQuAD question answering task by using high-dimensional low-precision embeddings in place of lower dimensional high-precision embeddings.

- In Appendix E.4 we present extended results comparing the eigenspace overlap scores of the different compression methods, for different compression rates and embeddings types. We show that uniform quantization can attain comparable or higher eigenspace overlap scores relative to the other compression methods, helping to explain the strong empirical performance of this compression method.
- In Appendix E.5 we present extended results on the correlations between downstream performance and the different measures of compression quality. We show that across the question answering, sentiment analysis, and GLUE tasks we consider, the eigenspace overlap score consistently attains higher Spearman correlation with downstream performance than the other measures of compression quality (PIP loss,  $\Delta$ ,  $\Delta_{\max}$ ).
- In Appendix E.6 we show that the eigenspace overlap score also correlates better with downstream performance than the  $\Delta_1$  and  $\Delta_2$  compression quality metrics, across a range of tasks.
- In Appendix E.7 we show that our claim that the eigenspace overlap score correlates better with downstream performance than  $\Delta$  and  $\Delta_{\max}$  is robust to the choice of the parameter  $\lambda$  used when computing the values of  $\Delta$  and  $\Delta_{\max}$ .
- In Appendix E.8 we show that the eigenspace overlap score is a more robust selection criterion for choosing between pairs of compressed embeddings than the other measures of compression quality.
- In Appendix E.9 we compare the downstream performance of embeddings compressed using deterministic vs. stochastic uniform quantization. We show these methods perform similarly, though the deterministic quantization performs slightly better at precision  $b = 1$ .

We present all these results in more detail below.

## E.1 Downstream Performance vs. Compression Rate: Pre-Trained Embeddings

In Figures 6 (GloVE), 7 (fastText), and 8 (BERT), we show the downstream performance of the embeddings compressed using different compression methods, across question answering, sentiment analysis, and GLUE tasks. We show that the simple uniform quantization method can match or outperform the other compression methods across these tasks. We also observe that for the GLUE tasks (Figure 8), freezing the WordPiece embeddings during the BERT model fine-tuning does not observably hurt downstream performance.

## E.2 Downstream Performance vs. Compression Rate: Task-Specific Embeddings

The main focus of our work is on understanding the downstream performance of NLP models trained using compressed pre-trained word embeddings. Recently, Khrulkov et al. [16] proposed compressing word embedding matrices by parameterizing them as a product of tensors, and then learning the entries of these tensors jointly with the downstream NLP model in a task-specific, end-to-end fashion; they call this method a *Tensor Train (TT) decomposition* of the embedding matrix. In this section, we show that we can apply uniform quantization to compressing task-specific word embeddings, and attain competitive downstream performance with the TT method.

**Task details** We consider the IWSLT’14 German-to-English translation task [4]. We use a six-layer Transformer [36] based translation model for this task, and use the Fairseq [25] implementation of this model. In our experiments, across all compression rates and compression methods, we train for 50000 steps, and use the same model size with a 512-dimensional transformer hidden layer; thus, the uncompressed embeddings are 512 dimensional. We use the default training and inference hyperparameters for this German-to-English translation task in the Fairseq repository; we list the values of these hyperparameters in Table 12. To be compatible with the Fairseq implementation, we run these experiments using PyTorch 1.0.

**Compression method details** We now provide details on how we apply the different compression methods in this task-specific setting. Note that because TT can achieve compression rates greater than  $32\times$ , we run experiments both above and below this compression rate.

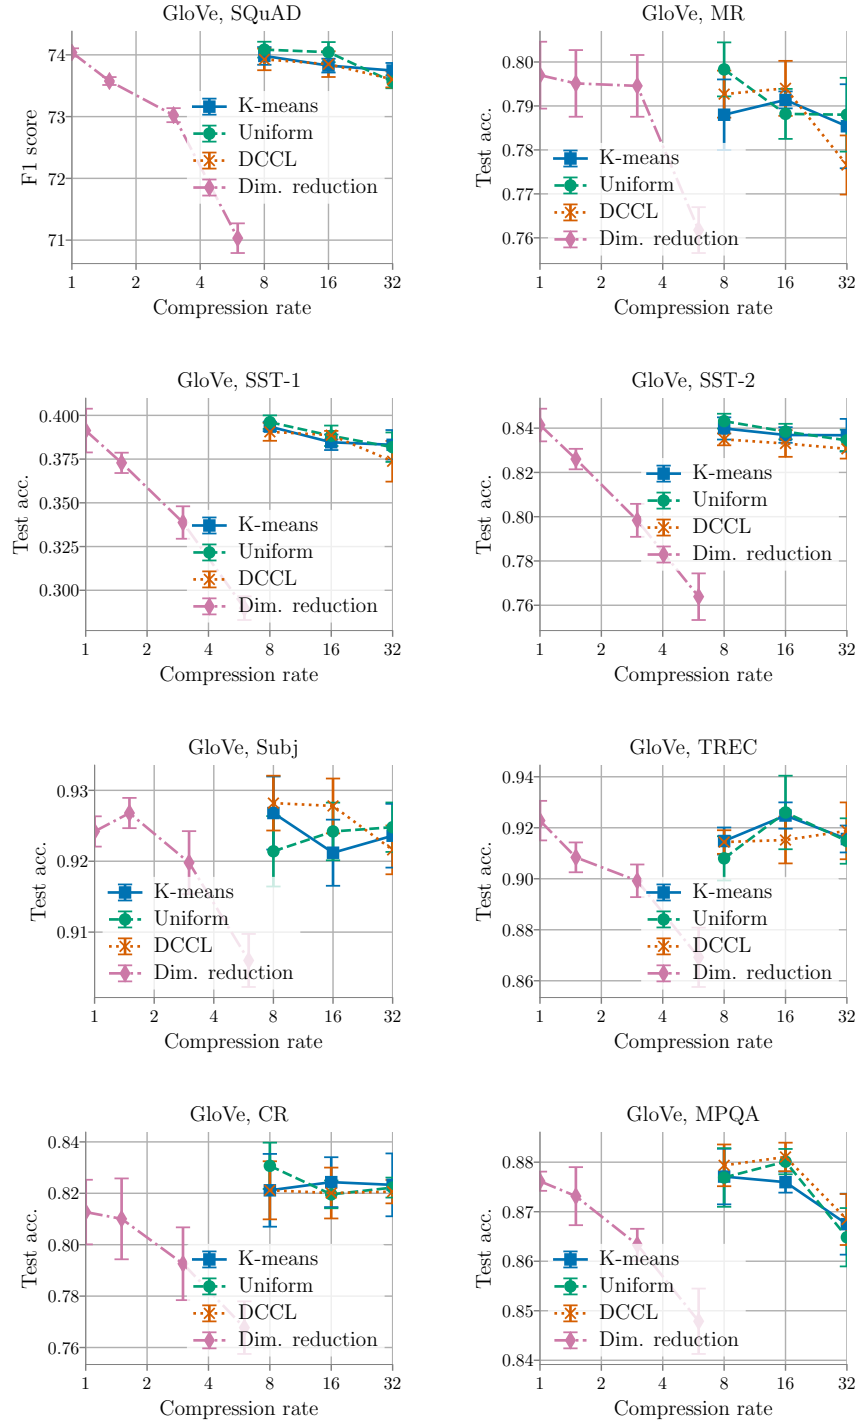


Figure 6: **Downstream performance vs. compression rate for compressed GloVe embeddings.** We evaluate the downstream performance of the different compression methods on question answering and sentiment analysis tasks, across different compression rates. For question answering, we use the SQuAD dataset, and for sentiment analysis we use the MR, SST-1, SST-2, Subj, TREC, CR and MPQA datasets. We show average performance across five random seeds, with error bars indicating standard deviations.

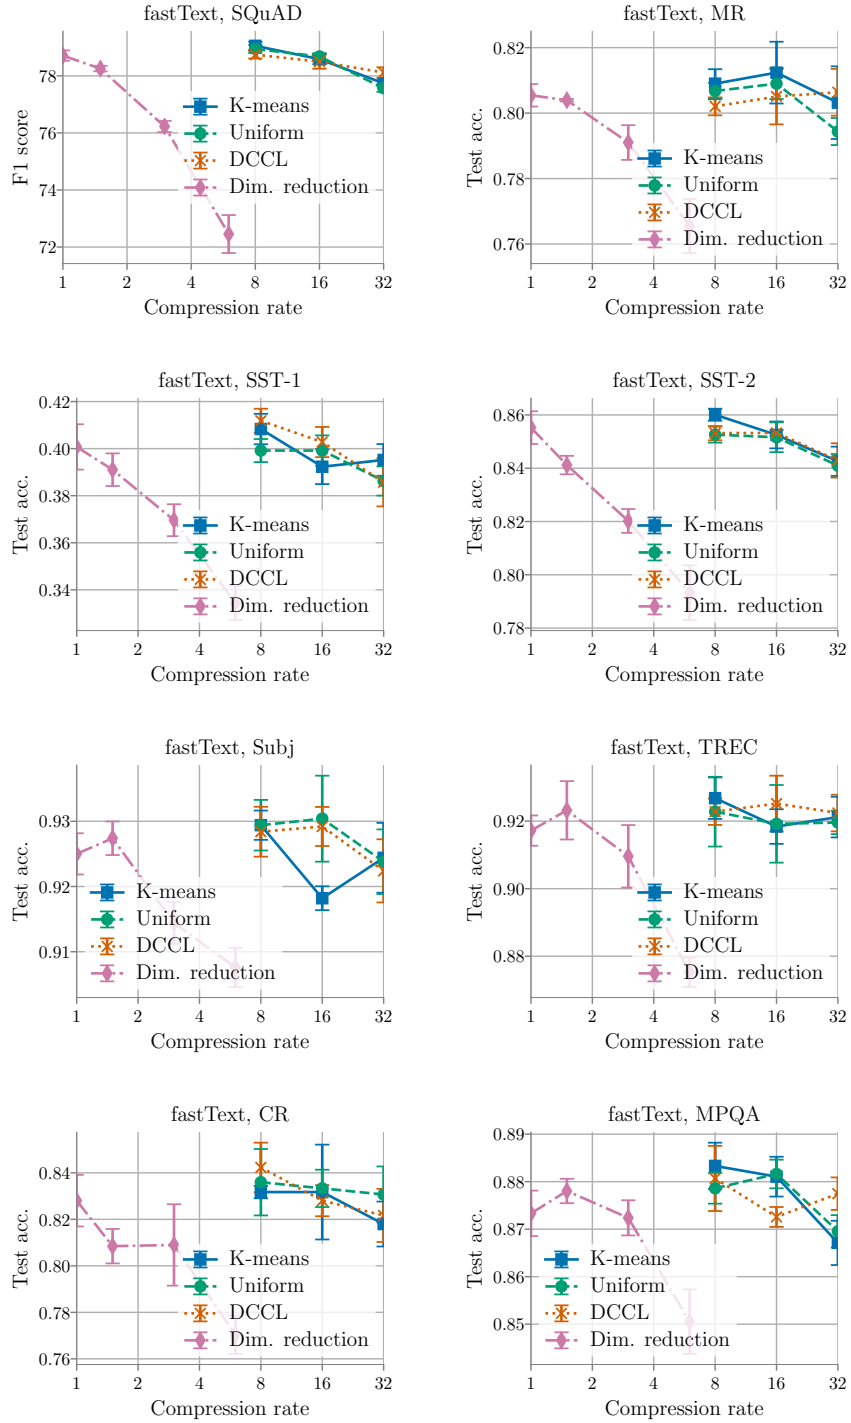


Figure 7: **Downstream performance vs. compression rate for compressed fastText embeddings.** We evaluate the downstream performance of the different compression methods on question answering and sentiment analysis tasks, across different compression rates. For question answering, we use the SQuAD dataset, and for sentiment analysis we use the MR, SST-1, SST-2, Subj, TREC, CR and MPQA datasets. We show average performance across five random seeds, with error bars indicating standard deviations.

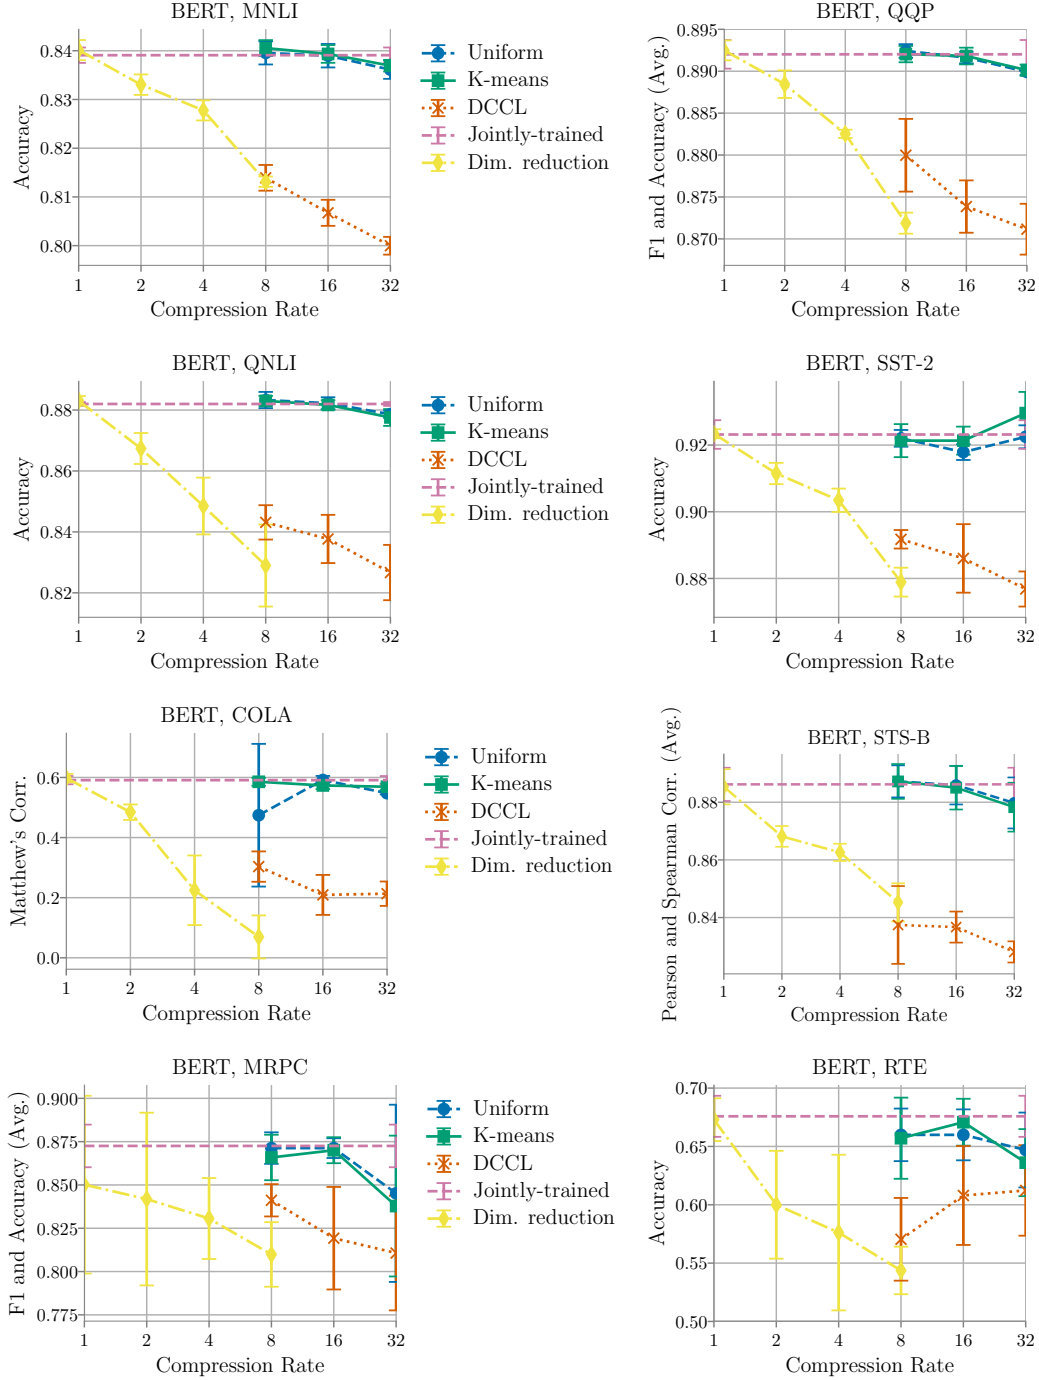


Figure 8: **Downstream performance vs. compression rate for compressed BERT WordPiece embeddings.** We evaluate the downstream performance of the different compression methods on all GLUE tasks except WNLI (as discussed in Appendix D.1), across different compression rates. In these plots, the horizontal dashed pink line marks the performance of the BERT model fine-tuned with uncompressed and unfrozen WordPiece embeddings. We show average performance across five random seeds, with error bars indicating standard deviations.



Table 11: The optimal learning rates  $\eta$  and dictionary sizes  $k$  for DCCL for compressing the task-specific embeddings for the IWSLT’14 translation task.

Compression rate	8×	16×	32×	64×	128×	256×
$k$	16	16	4	4	4	2
$\eta$	0.0003	0.0003	0.001	0.001	0.001	0.001

- **Dimensionality reduction:** We randomly initialize lower-dimensional embeddings, and train the parameters of these embeddings jointly with the rest of the model.
- **Uniform quantization:** We jointly train the full-precision embedding matrix and the transformer model for the first half of the training steps; we then compress this embedding matrix with uniform quantization (Algorithm 1), and keep the embedding parameters fixed for the remainder of training. To attain a compression rate  $c > 32$ , we perform the first half of training using lower-dimensional embeddings (compression rate  $c/32$ ), and then apply uniform quantization to these lower-dimensional embeddings with compression rate 32.
- **K-means:** We use the same protocol as we do for uniform quantization, but apply the k-means compression method in place of uniform quantization.
- **DCCL:** As we do for uniform quantization and k-means, we jointly train the full-precision embedding matrix and the transformer model for the first half of the training steps; we then compress the embeddings with DCCL, and perform the rest of training with the embedding parameters fixed. We grid search the dictionary size  $k \in \{2, 4, 8, 16, 32, 64\}$  and the learning rate  $\eta \in \{0.00003, 0.0001, 0.0003, 0.001, 0.003, 0.01\}$  for DCCL, and pick the combination of values which minimizes the embedding reconstruction error with respect to the embeddings generated in the first half of training. We show the optimal hyperparameters for each compression rate in Table 11.
- **Tensor Train:** We use the TT method in the manner described in the original paper [16]. For each compression rate, there are two hyperparameters that must be tuned—the number of tensor factors and the “TT-rank” of these factors. We consider 3 and 4 as the number of factors, following the values used in the paper [16], and pick the one which gives the lowest validation perplexity. Given the number of factors, the TT-rank of these factors is automatically determined for a given compression rate.

**Results** In Figure 9 we plot the average test BLEU4 score across five random seeds for the compression methods described above, at a wide range of compression rates; because for some random seeds the TT method attains very low BLEU scores, for the TT method we plot the BLEU4 score of the seed which performs *best*. We observe that the uniform quantization and k-means methods generally achieve better BLEU4 score than the TT method up to compression rate 128×, and that the dimensionality reduction method performs significantly worse than the other methods beyond compression rate 8×. These observations suggest that uniform quantization and k-means can be effectively applied to compress task-specific embeddings.

### E.3 Dimension vs. Precision Trade-Off

We show that in the memory constrained setting, using low-precision high-dimensional embeddings typically outperforms using high-precision low-dimensional embeddings which occupy the same memory. To demonstrate this, we train GloVe embeddings (details below) of dimensions  $d \in \{25, 50, 100, 200, 400\}$ , and then compress each of these embeddings using uniform quantization with precisions  $b \in \{1, 2, 4, 8, 16, 32\}$  (32 bits represents no compression). We then train DrQA models [5] using all of these embeddings on the SQuAD dataset [32], and CNN models [18] on the SST-1 sentiment analysis dataset. In Figure 10 we present the downstream performance of all of these models ( $y$ -axis) in terms of the memory occupied by the embeddings ( $x$ -axis). As we can see, across a range of memory budgets, it is optimal to use low-precision (1 bit) high-dimensional embeddings, as this allows for using the largest dimension possible under that memory budget.

**GloVe embedding training details** We train GloVe embeddings on a full English Wikimedia dump on December 4, 2017 which was pre-processed by a fastText script<sup>23</sup> while keeping the letter

<sup>23</sup><https://github.com/facebookresearch/fastText/blob/master/get-wikimedia.sh>

Table 12: The hyperparameters we use for our experiments on the IWSLT’14 German-to-English translation task.

Hyperparameter	Value
Optimizer	Adam
Adam decay rates for 1st moment $\beta_1$	0.9
Adam decay rates for 2nd moment $\beta_2$	0.999
Adam $\epsilon$	$10^{-8}$
Training steps	50000
Learning rate schedule	$10^{-7} + (5 * 10^{-4} - 10^{-7})n/4000$ for step $n \leq 4000$ $5 * 10^{-4} * \sqrt{4000/n}$ for step $n > 4000$
Warmup initial learning rate	$10^{-7}$
Dropout	0.3
Weight decay	0.0001
Beam search width	5
Transformer hidden dimension	512




Figure 9: **Downstream performance vs. compression rate: task-specific embeddings.** We plot the average BLEU4 test performance for compressed task-specific embeddings on the IWSLT’14 German-to-English translation task across five random seeds (standard deviations indicated with error bars). note that because for some random seeds the TT method attains very low BLEU scores, in this plot we report the *best* performance for the TT method across the five random seeds. We observe that the uniform quantization and k-means compression methods generally achieve better BLEU4 test performance than the TT method for compression rates up to 128 $\times$ .

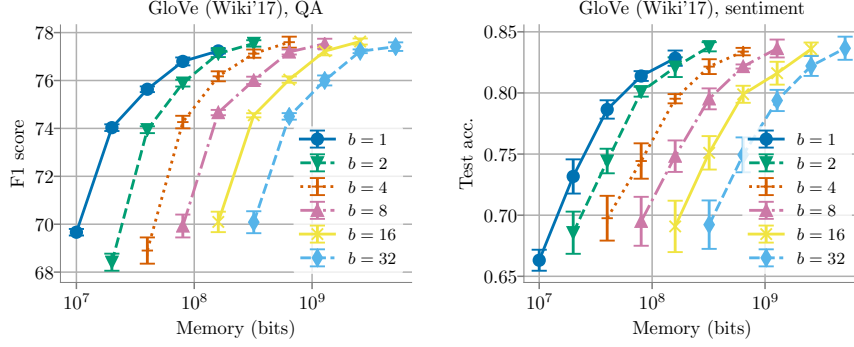


Figure 10: **Dimension vs. precision trade-off.** We plot the downstream performance on question answering (SQuAD, left) and sentiment analysis (SST-2, right) of GloVe embeddings of dimensions  $d \in \{25, 50, 100, 200, 400\}$  compressed with uniform quantization with precisions  $b \in \{1, 2, 4, 8, 16, 32\}$ , as a function of the memory occupied by the embeddings. We observe that embeddings compressed with 1-bit precision typically demonstrate the best downstream performance under a range of memory budgets. We show average performance across five random seeds, with error bars indicating standard deviations.




Figure 11: **Eigenspace overlap score vs. compression rate.** We plot the eigenspace overlap scores attained by different compression methods at different compression rates, for GloVe, fastText, and BERT WordPiece embeddings. Across embedding types and compression rates, we observe that uniform quantization attains similar or higher eigenspace overlap scores than the other compression methods. We show average eigenspace overlap scores across five random seeds, with error bars indicating standard deviations.

cases and digits. We use the GloVe Github repository<sup>24</sup> for embedding training. We use a vocabulary size of 400000, a window size of 15, a learning rate of 0.05, and train for 50 epochs.

#### E.4 Eigenspace Overlap Score vs. Compression Rate

In Figure 11, we plot the eigenspace overlap scores attained by the different compression methods at different compression rates for GloVe, fastText, and BERT WordPiece embeddings. We observe that uniform quantization consistently attains higher or matching eigenspace overlap scores than the other compression methods. Based on the theoretical connection between the eigenspace overlap score and downstream performance, this empirical observation helps explain the strong downstream performance of embeddings compressed with uniform quantization.

#### E.5 Downstream Performance vs. Measures of Compression Quality

We show across tasks and embedding types that the eigenspace overlap score correlates better with downstream performance than the other measures of compression quality. In Figures 12, 13,

<sup>24</sup><https://github.com/stanfordnlp/GloVe>

Table 13: **Spearman correlations  $\rho$  between compression quality measures and sentiment analysis performance.** Within each entry of the table, the correlations are presented in terms of ‘GloVe  $|\rho|$  / fastText  $|\rho|$ .’ Note that we present the absolute values of the correlation coefficients, with higher absolute values indicating stronger correlation.

	MR	SST-1	SST-2	Subj	TREC	CR	MPQA
PIP loss	0.36/0.03	0.46/0.25	0.29/0.21	0.30/0.13	0.14/0.05	0.10/0.03	0.49/0.18
$\Delta$	0.29/0.39	0.33/0.29	0.39/0.40	0.26/0.16	0.33/0.29	0.11/0.34	0.41/0.40
$\Delta_{\max}$	<b>0.39</b> /0.41	0.51/0.60	0.41/0.62	<b>0.32</b> /0.49	0.23/0.30	0.12/0.40	<b>0.60</b> /0.39
$1 - \mathcal{E}$	0.29/ <b>0.59</b>	<b>0.75</b> / <b>0.73</b>	<b>0.72</b> / <b>0.83</b>	0.27/ <b>0.58</b>	<b>0.49</b> / <b>0.32</b>	<b>0.40</b> / <b>0.55</b>	<b>0.60</b> / <b>0.55</b>

Table 14: **Spearman correlations  $\rho$  between compression quality measures and GLUE task performance.** Note that we present the absolute values of the correlation coefficients, with higher absolute values indicating stronger correlation.

	MNLI	QQP	QNLI	SST-2	CoLA	STS-B	MRPC	RTE
PIP loss	0.45	0.45	0.43	0.18	0.32	0.41	0.28	0.22
$\Delta$	0.44	0.36	0.36	0.25	0.25	0.43	0.23	0.12
$\Delta_{\max}$	0.86	0.86	0.85	0.67	0.75	0.84	0.59	0.58
$1 - \mathcal{E}$	<b>0.92</b>	<b>0.93</b>	<b>0.92</b>	<b>0.83</b>	<b>0.86</b>	<b>0.87</b>	<b>0.62</b>	<b>0.66</b>

and 14, we plot the downstream performance ( $y$ -axis) of the compressed GloVe, fastText, and BERT WordPiece embeddings (respectively) on a variety of tasks, as a function of the different measures of compression quality ( $x$ -axis). For GloVe and fastText, we show performance on question answering (SQuAD) and on the largest sentiment analysis dataset (SST-1). For BERT, we show performance on MNLI and QQP, the two largest GLUE datasets. We see in these plots that the eigenspace overlap score generally aligns quite well with downstream performance, while the other measures of compression quality often do not. To quantify this observation, we measure the Spearman correlations between the downstream performances of the embeddings we compressed, and the various measures of compression quality. We include these correlations for all the sentiment analysis tasks for the GloVe and fastText embeddings in Table 13, and for all the GLUE tasks for the BERT WordPiece embeddings in Table 14. From these results, we can see that across different tasks and embedding types, the eigenspace overlap score generally correlates better with downstream performance than the other measure of compression quality.

## E.6 Downstream Performance vs. $1/(1 - \Delta_1)$ and $\Delta_2$

We show two main results: First, we show examples of compressed embeddings that have large values of  $\frac{1}{1 - \Delta_1}$  or  $\Delta_2$ , but which still attain strong downstream performance; because large values of  $\frac{1}{1 - \Delta_1}$  or  $\Delta_2$  imply large worst-case bounds on the generalization error of the embeddings [41], these observations demonstrate that the worst-case bounds are too loose to explain the empirical results. Second, we show that the eigenspace overlap score generally attains stronger correlation with downstream performance than both  $1/(1 - \Delta_1)$  and  $\Delta_2$ .

For the first result, we can see in Figure 15 that there are points with large  $\frac{1}{1 - \Delta_1}$ , for example, but where the downstream performance is still quite close to the full-precision embedding performance. For the second result, we show in Table 15 that the eigenspace overlap score attains higher Spearman correlation with downstream performance than  $1/(1 - \Delta_1)$  and  $\Delta_2$  across a range of tasks.

## E.7 Downstream Performance vs. $\Delta_{\max}$ and $\Delta$ with different $\lambda$ values

In Section 4, we showed across numerous tasks and embedding types that the eigenspace overlap score typically attains stronger correlation with downstream performance than the other measures of compression quality, including  $\Delta_{\max}$  and  $\Delta$ . For these results, we computed  $\Delta_{\max}$  and  $\Delta$  with the parameter  $\lambda$  being the smallest non-zero eigenvalue of the Gram matrix of the uncompressed embeddings (see Section 2.2 for a review of how  $\lambda$  is used when calculating these measures). We now show these results are robust to the choice of  $\lambda$ . Specifically, in Table 16 we show the Spearman correlations attained by  $\Delta_{\max}$  and  $\Delta$  with different  $\lambda$  values. Letting  $\lambda_{\min}$  and  $\lambda_{\max}$  be

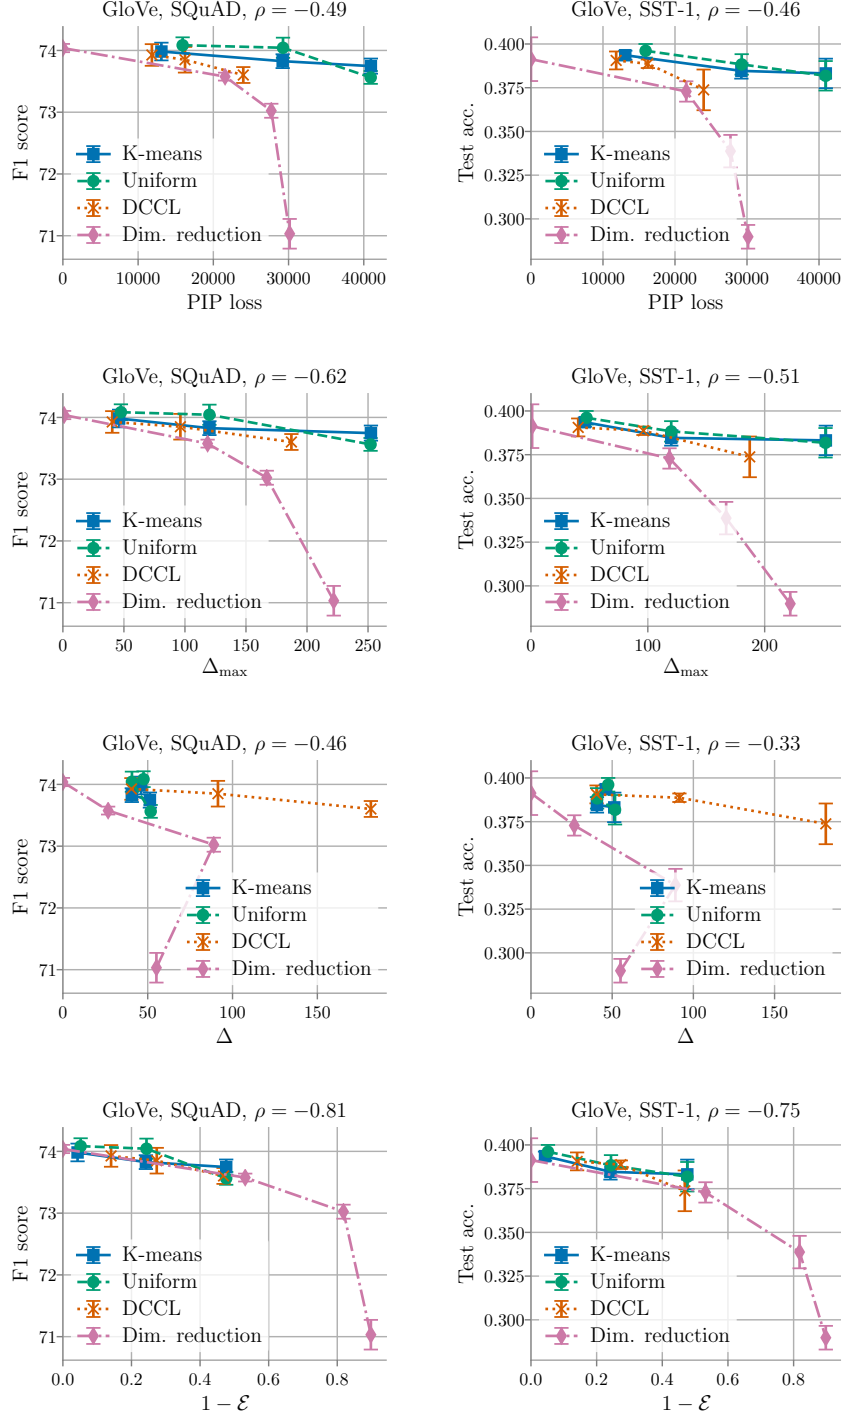


Figure 12: **Downstream performance vs. measures of compression quality (GloVe embeddings).** We plot the performance of compressed GloVe embeddings on question answering (SQuAD, left column) and sentiment analysis (SST-1, right column), in terms of the different measures of compression quality for these embeddings. We can see that the eigenspace overlap score  $\mathcal{E}$  generally aligns better with downstream performance than the other measures of compression quality. To quantify this, in the title of each plot we include the Spearman correlation  $\rho$  between downstream performance and the measure of compression quality for that plot. We can see that the eigenspace overlap score attains the strongest correlations with downstream performance, as it has the largest values for  $|\rho|$ .

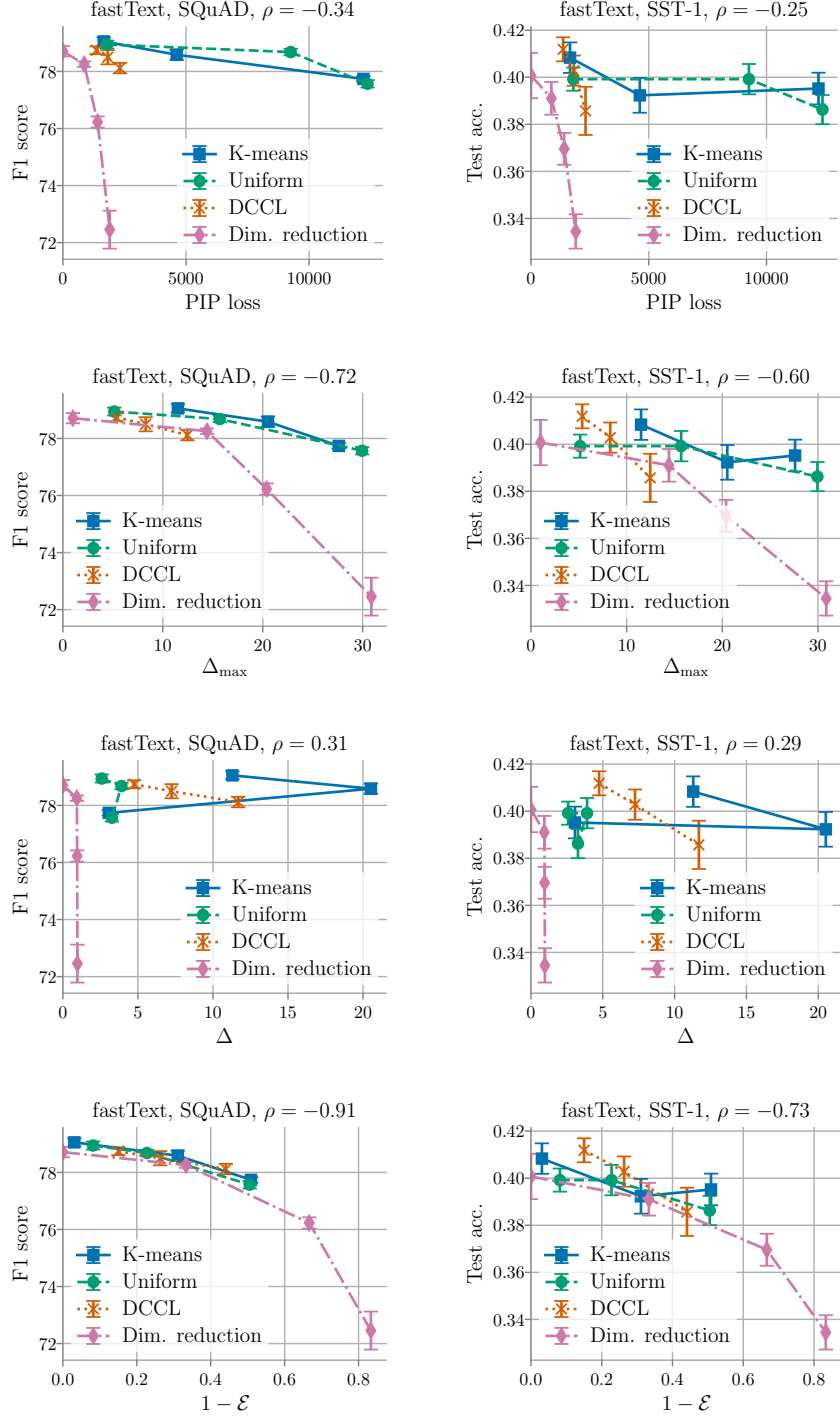


Figure 13: **Downstream performance vs. measures of compression quality (fastText embeddings).** We plot the performance of compressed fastText embeddings on question answering (SQuAD, left column) and sentiment analysis (SST-1, right column), in terms of the different measures of compression quality for these embeddings. We can see that the eigenspace overlap score  $\epsilon$  generally aligns better with downstream performance than the other measures of compression quality. To quantify this, in the title of each plot we include the Spearman correlation  $\rho$  between downstream performance and the measure of compression quality for that plot. We can see that the eigenspace overlap score attains the strongest correlations with downstream performance, as it has the largest values for  $|\rho|$ .

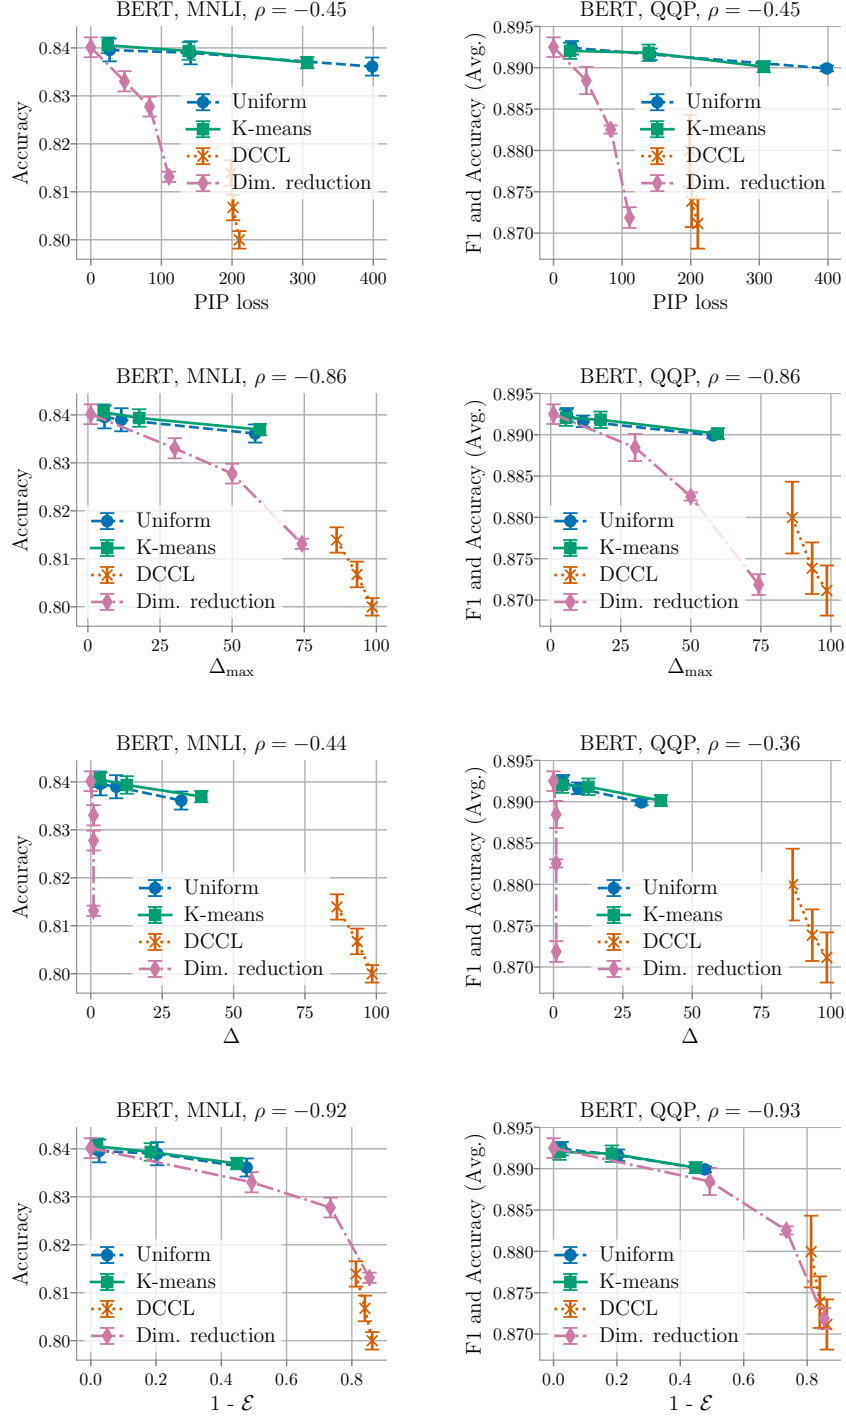


Figure 14: **Downstream performance vs. measures of compression quality (BERT WordPiece embeddings).** We plot the performance of compressed BERT WordPiece embeddings on the two largest GLUE datasets (MNLI, left column; QQP, right column) in terms of the different measures of compression quality for these embeddings. We can see that the eigenspace overlap score  $\mathcal{E}$  generally aligns better with downstream performance than the other measures of compression quality. To quantify this, in the title of each plot we include the Spearman correlation  $\rho$  between downstream performance and the measure of compression quality for that plot. We can see that the eigenspace overlap score attains the strongest correlations with downstream performance, as it has the largest values for  $|\rho|$ .

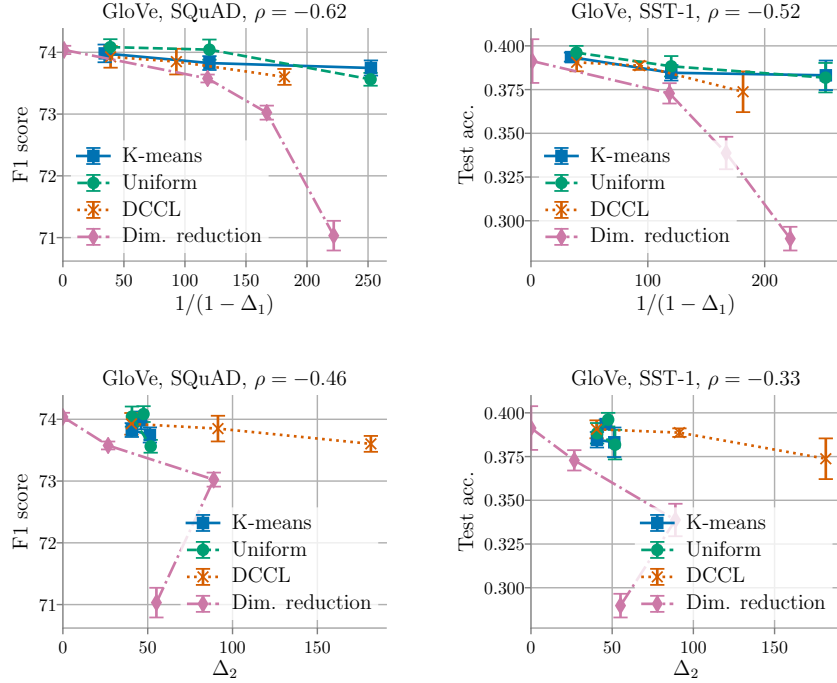


Figure 15: **Downstream performance vs.  $1/(1 - \Delta_1)$  and  $\Delta_2$  (GloVe embeddings).** We plot the performance of compressed GloVe embeddings on question answering (SQuAD, left column) and sentiment analysis (SST-1, right column), in terms of the  $1/(1 - \Delta_1)$  and  $\Delta_2$  measures of compression quality. We can see that there are compressed embeddings with large values  $\frac{1}{1-\Delta_1}$ , but where the downstream performance is still quite close to the full-precision embedding performance. Additionally, we can see that from a visual perspective, these compression quality measures do not align very well with downstream performance. For example, the dimensionality reduction embeddings with compression ratio  $6\times$  attain smaller  $1/(1 - \Delta_1)$ , but worse F1 score on SQuAD, than uniform quantization embeddings with compression ratio  $32\times$ .

Table 15: **Spearman correlations between compression quality measures and downstream performance.** On the SQuAD (question answering), SST-1 (sentiment analysis), MNLI (natural language inference), and QQP (question pair matching) tasks, the eigenspace overlap score  $\mathcal{E}$  attains higher Spearman correlation (absolute value) with downstream performance than  $1/(1 - \Delta_1)$  and  $\Delta_2$ .

Dataset	SQuAD		SST-1		MNLI	QQP
Embedding	GloVe	fastText	GloVe	fastText	BERT WordPiece	BERT WordPiece
$1/(1 - \Delta_1)$	0.62	0.80	0.52	0.65	0.87	0.87
$\Delta_2$	0.46	0.48	0.33	0.44	0.30	0.20
$1 - \mathcal{E}$	<b>0.81</b>	<b>0.91</b>	<b>0.75</b>	<b>0.73</b>	<b>0.92</b>	<b>0.93</b>



Table 16: **Spearman correlations between  $\Delta$  and  $\Delta_{\max}$  and downstream performance, for different  $\lambda$  values.** On the SQuAD (question answering), SST-1 (sentiment analysis), MNLI (natural language inference), and QQP (question pair matching) tasks, the eigenspace overlap score  $\mathcal{E}$  attains higher Spearman correlation (absolute value) with downstream performance than  $\Delta_{\max}$  and  $\Delta$  computed with different  $\lambda$  values.

Dataset		SQuAD		SST-1		MNLI	QQP
Embedding		GloVe	fastText	GloVe	fastText	BERT WordPiece	BERT WordPiece
$\Delta_{\max},$	$\lambda = \lambda_{\min}/100$	0.66	0.71	0.54	0.63	0.47	0.56
$\Delta_{\max},$	$\lambda = \lambda_{\min}/10$	0.65	0.73	0.54	0.61	0.47	0.57
$\Delta_{\max},$	$\lambda = \lambda_{\min}$	0.62	0.72	0.51	0.60	0.38	0.56
$\Delta_{\max},$	$\lambda = \lambda_{\min} \times 10$	0.61	0.65	0.53	0.51	0.35	0.43
$\Delta_{\max},$	$\lambda = \lambda_{\min} \times 100$	0.25	0.49	0.18	0.43	0.13	0.36
$\Delta_{\max},$	$\lambda = \lambda_{\max}$	0.15	0.08	0.22	0.03	0.49	0.08
$\Delta,$	$\lambda = \lambda_{\min}/100$	0.41	0.31	0.27	0.30	0.51	0.05
$\Delta,$	$\lambda = \lambda_{\min}/10$	0.41	0.32	0.27	0.30	0.51	0.05
$\Delta,$	$\lambda = \lambda_{\min}$	0.46	0.31	0.33	0.29	0.57	0.04
$\Delta,$	$\lambda = \lambda_{\min} \times 10$	0.42	0.00	0.28	0.05	0.55	0.28
$\Delta,$	$\lambda = \lambda_{\min} \times 100$	0.70	0.32	0.60	0.27	0.87	0.26
$\Delta,$	$\lambda = \lambda_{\max}$	0.35	0.01	0.31	0.03	0.10	0.02
$1 - \mathcal{E}$		<b>0.81</b>	<b>0.91</b>	<b>0.75</b>	<b>0.73</b>	<b>0.92</b>	<b>0.93</b>

Table 17: **The robustness of each measure of compression quality as a selection criterion.** Across all pairs of compressed embeddings from our experiments, we measure for each task the maximum difference in performance between the embedding selected by each measure of compression quality and the one which performs best on the task. We report these results in the table below, and observe that the eigenspace overlap score  $\mathcal{E}$  attains lower maximum performance differences than the other measures of compression quality.

Dataset	SQuAD		SST-1		MNLI	QQP
Embedding	GloVe	fastText	GloVe	fastText	BERT WordPiece	BERT WordPiece
PIP loss	0.03	0.08	0.11	0.08	0.04	0.02
$\Delta_{\max}$	0.03	0.03	0.11	0.05	0.02	0.02
$\Delta$	0.03	0.08	0.11	0.09	0.03	0.02
$1 - \mathcal{E}$	<b>0.01</b>	<b>0.01</b>	<b>0.04</b>	<b>0.03</b>	<b>0.01</b>	<b>0.01</b>

the smallest and largest eigenvalues of the uncompressed embedding Gram matrix, we consider  $\lambda \in \{\lambda_{\min}/100, \lambda_{\min}/10, \lambda_{\min}, \lambda_{\min} \times 10, \lambda_{\min} \times 100, \lambda_{\max}\}$  for this table. We observe that the eigenspace overlap score attains stronger correlation with downstream performance across the tasks and embedding types in this table than  $\Delta$  and  $\Delta_{\max}$ , across all the  $\lambda$  values listed above.

## E.8 The Robustness of the Measures of Compression Quality as Selection Criteria

In Section 4.3 we argued that the eigenspace overlap score is a more accurate and robust selection criterion for choosing between compressed embeddings than the other measures of compression quality. We showed in Table 2 the selection error rates attained by the various measures of compression quality across different tasks and embeddings types. Here we provide detailed results on the robustness of the various measures of compression quality when used as selection criteria. To quantify the robustness of each measure of compression quality as a selection criterion, we measure for each task the maximum difference in performance, across all pairs of compressed embeddings from our experiments, between the embedding which performs best and the one which is selected by the measure of compression quality. We report these results in Table 17 for GloVe and fastText embeddings on the question answering (SQuAD) and sentiment analysis (SST-1) tasks, and for BERT WordPiece embeddings on the language inference (MNLI) and question pair classification (QQP) tasks. We observe that the eigenspace overlap score can attain  $1.1 \times$  to  $5.5 \times$  lower maximum performance differences than the next best measures of compression quality.

### E.9 Stochastic vs. Deterministic Uniform Quantization

Thus far, all the uniform quantization experiments we have presented on question answering, sentiment analysis, and GLUE tasks have used deterministic rounding. However, our theoretical analysis on the expected eigenspace overlap score of uniformly quantized embeddings assumed unbiased stochastic quantization is used. In this section, we show that (1) stochastic and deterministic uniform quantization perform similarly on downstream tasks, and that (2) the eigenspace overlap score still correlates well with downstream performance when using stochastic quantization instead of deterministic quantization. In Figure 16, we compare the downstream performance of deterministic and stochastic quantization on the SQuAD question answering task and on the SST-1 sentiment analysis task. We can observe that uniform and deterministic quantization perform similarly, although at 1-bit precision deterministic quantization performs slightly better than stochastic quantization. We then show in Figure 17 that regardless of whether we use deterministic or stochastic quantization, the eigenspace overlap score correlates better with downstream performance across compression methods than the other measures of compression quality.

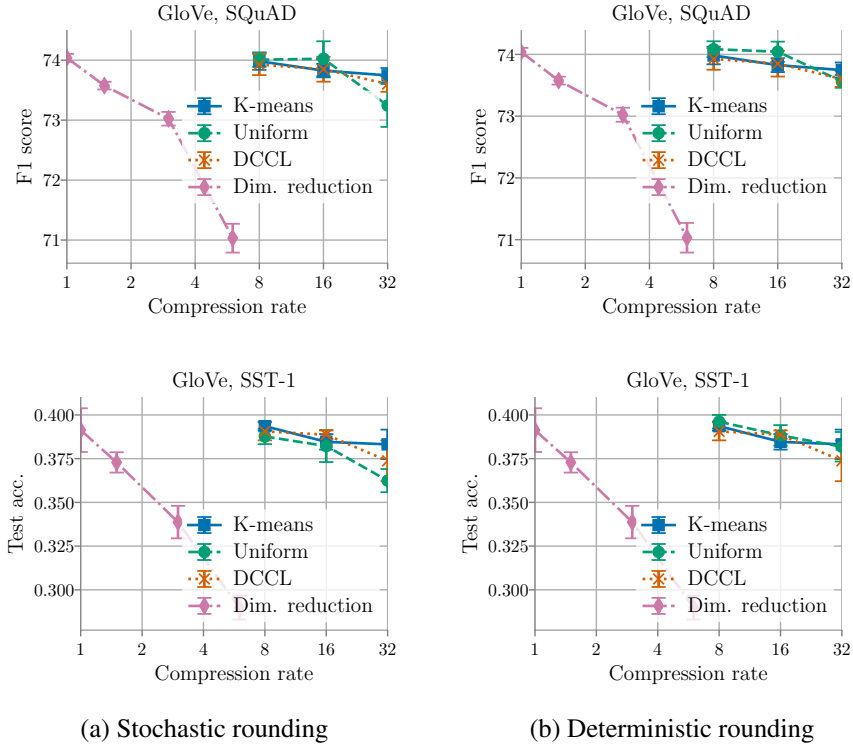


Figure 16: **Downstream performance vs. compression rate for deterministic vs. stochastic uniform quantization.** We can observe, using compressed GloVe embeddings on both the SQuAD question answering task and the SST-1 sentiment analysis task, that stochastic uniform quantization (left plots) performs similarly to deterministic uniform quantization (right plots).

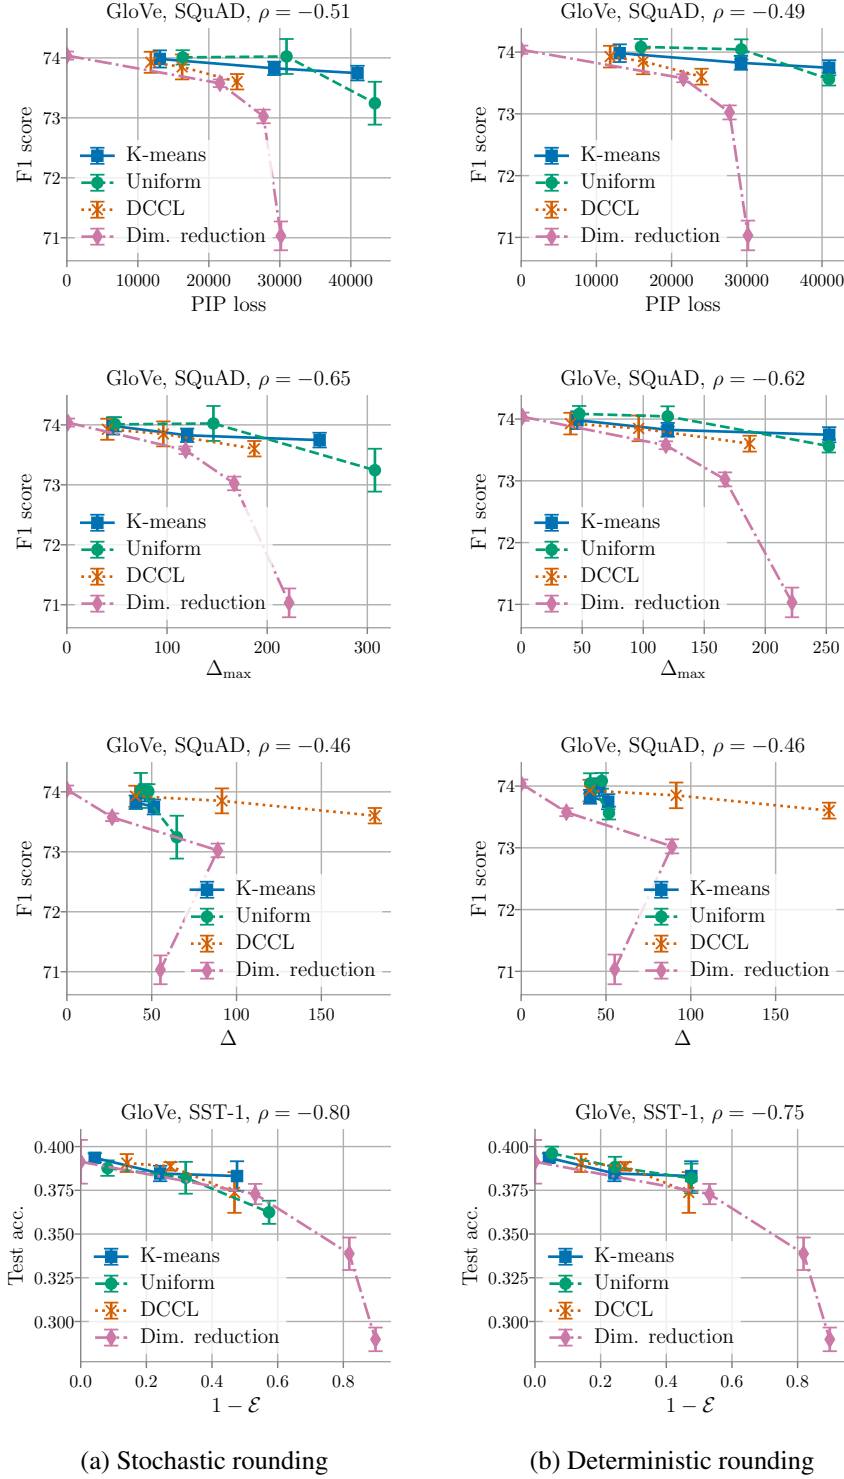


Figure 17: **Downstream performance vs. measures of compression quality for deterministic vs. stochastic uniform quantization.** We can see that regardless of whether stochastic (left plots) or deterministic (right plots) quantization is used, the eigenspace overlap score correlates better with downstream performance than the other measures of compression quality (as quantified by the Spearman correlations  $\rho$  in the plot titles).