

Matching-oriented Product Quantization For Ad-hoc Retrieval

Shitao Xiao^{1†}, Zheng Liu^{2*}, Yingxia Shao^{1*}, Defu Lian³, Xing Xie²

1: Beijing University of Posts and Telecommunications, Beijing, China

2: Microsoft Research Asia, Beijing, China

3: University of Science and Technology of China, Hefei, China

{stxiao, shaoyx}@bupt.edu.cn

{zhengliu, xingx}@microsoft.com

defulian@ustc.edu.cn

Abstract

Product quantization (PQ) is a widely used technique for ad-hoc retrieval. Recent studies propose supervised PQ, where the embedding and quantization models can be jointly trained with supervised learning. However, there is a lack of appropriate formulation of the joint training objective; thus, the improvements over previous non-supervised baselines are limited in reality. In this work, we propose the Matching-oriented Product Quantization (**MoPQ**), where a novel objective Multi-noulli Contrastive Loss (**MCL**) is formulated. With the minimization of MCL, we are able to maximize the matching probability of query and ground-truth key, which contributes to the optimal retrieval accuracy. Given that the exact computation of MCL is intractable due to the demand of vast contrastive samples, we further propose the Differentiable Cross-device Sampling (**DCS**), which significantly augments the contrastive samples for precise approximation of MCL. We conduct extensive experimental studies on four real-world datasets, whose results verify the effectiveness of MoPQ. The code is available at <https://github.com/microsoft/MoPQ>.

1 Introduction

Ad-hoc retrieval is critical for many intelligent services, like web search and recommender systems (Liu et al., 2021; Zhu et al., 2021; Ma et al., 2019). Given a query (e.g., a search request from user), ad-hoc retrieval selects relevant keys (e.g., webpages) from a massive set of candidates. It is usually treated as an approximate nearest neighbour search (ANNS) problem, where product quantization (PQ) (Jegou et al., 2010) is one of the most popular solutions thanks to its competitive memory and time efficiency. PQ is built upon

“codebooks”, with which an input embedding can be quantized into a Cartesian product of “codewords” (preliminaries about the codebooks and codewords will be given in Section 2.1). By this means, the original embedding can be compressed into a highly compact representation. More importantly, it significantly improves the retrieval efficiency, as query and key’s similarity can be approximated based on the pre-computed distances between query and codewords.

Existing Limitation. The original PQ algorithms (Jegou et al., 2010; Ge et al., 2013) are *non-supervised*: based on the well-trained embeddings, the quantization model is learned with heuristic algorithms (e.g., k -means). In recent years, many works are dedicated to *supervised PQ* (Cao et al., 2016, 2017; Klein and Wolf, 2019; Gao et al., 2019; Chen et al., 2020), where the embedding model and the quantization model are trained jointly. The supervised PQ requires an explicit training objective for the quantization model. Most of the time, “the minimization of **reconstruction loss**” is used for granted: the distortion between the original embedding (\mathbf{z}) and its quantization result ($\tilde{\mathbf{z}}$) needs to be reduced as much as possible for all the keys (k) in database: $\min \sum_k \|\mathbf{z}^k - \tilde{\mathbf{z}}^k\|_2$. The above objective is seemingly plausible, as a “small enough distortion” will make the quantized embeddings equally expressive as the original embeddings. However, it implicitly hypothesizes that the distortion can be made sufficiently small, which is not always realistic in practice. This is because the reconstruction loss is subject to a lower-bound determined by the codebooks scale. As over-scaled codebooks result in prohibitive memory and time costs, there will always exist a positive reconstruction loss in reality. In this case, it can be proved that the minimization of reconstruction loss doesn’t lead to the

†. Work is done during the internship at Microsoft.

*. Corresponding author.

optimal retrieval accuracy. It’s also empirically verified that the supervised PQ’s advantage over the non-supervised baselines are limited and not consistently positive when the reconstruction loss minimization is taken as the training objective (see Section 2.2 and 4 for detailed analysis).

Our Work. To address the above challenge, we propose the Matching-oriented Product Quantization (**MoPQ**), with a novel objective **MCL** formulated to optimize PQ’s retrieval accuracy, together with a sample augmentation strategy **DCS** to ensure the effective minimization of MCL.

- The Multinoulli Contrastive Loss (MCL) is formulated as the new quantization training objective. The PQ-based ad-hoc retrieval can be probabilistically modeled by a cascaded generative process: 1) select codewords for the input key, based on which the quantized key embedding is composited; 2) sample query from the Multinoulli distribution determined by the quantized key embedding. The negative of the generation likelihood is referred as the Multinoulli Contrastive Loss (MCL). *By minimizing MCL, the expected query-key matching probability will be optimized, which means the optimal retrieval accuracy.*

- The contrastive sample augmentation is designed to facilitate the minimization of MCL. The computation of MCL is intractable as it requires the normalization over vast contrastive samples (the quantized embeddings of all the keys). Thus, it has to be approximated by sampling whose effect is severely affected by sample size. In this work, we propose the Differentiable Cross-device Sampling (DCS), where a distributed embedding sharing mechanism is introduced for contrastive sample augmentation. As the gradients are stopped at the cross-device shared embeddings, we propose a technique called the “combination of Primary and Image Losses”, where the shared embeddings are made virtually differentiable to keep the model update free from distortions.

In summary, our work identifies a long-existing defect about the training objective of supervised PQ; meanwhile, a simple but effective remedy is proposed, which optimizes the expected retrieval accuracy of PQ. We make extensive experimental studies with four benchmark text retrieval datasets, where our proposed methods significantly outperform the SOTA supervised PQ baselines. Our code and datasets will be made public-available to facilitate the research progress in related areas.

2 Revisit of Supervised PQ

We start with the preliminaries of PQ’s application in ad-hoc retrieval. Then, we analyze the defect of taking the reconstruction loss minimization as supervised PQ’s training objective.

2.1 Preliminaries

- **Product Quantization (PQ).** PQ is a popular approach for learning quantized representations. It is based on the foundation of M codebooks \mathbf{C} : $\{\mathbf{C}_1, \dots, \mathbf{C}_M\}$; each \mathbf{C}_i consists of L codewords: $\{\mathbf{C}_{i1}, \dots, \mathbf{C}_{iL}\}$; each \mathbf{C}_{ij} is a vector whose dimension is d/M . Given an embedding \mathbf{z} (with dimension d), it is firstly sliced into to M sub-vectors $[\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_M]$; then the sub-vector \mathbf{z}_i is assigned to one of the codewords of codebook \mathbf{C}_i , whose ID is denoted by the one-hot vector \mathbf{b}_i . The assignment is made by the codeword selection function, which maps each sub-vector \mathbf{z}_i to its most relevant codeword, e.g., $\mathbf{b}_i = \text{one_hot}(\text{argmin}_k \|\mathbf{z}_i - \mathbf{C}_{i*}\|_2)$. As a result, the embedding \mathbf{z} is quantized into a collection of codes: $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_M\}$, where the embedding itself is approximated by the concatenation of the assigned codewords: $\tilde{\mathbf{z}} = [\mathbf{C}_1\mathbf{b}_1, \dots, \mathbf{C}_M\mathbf{b}_M]$.

Non-supervised PQ takes the well-trained embeddings as input, and learns the quantization model with heuristic algorithms, like k -means. In contrast, supervised PQ jointly learns the embedding and quantization models based on labeled data (the paired query and key). Specifically, it learns the query and key’s embeddings: \mathbf{z}^q and \mathbf{z}^k , such that the query and key’s relationship can be predicted based on their inner product $\langle \mathbf{z}^q, \mathbf{z}^k \rangle$. Furthermore, it learns the codebooks where the quantized embeddings may preserve the same expressiveness as the original embeddings.

- **PQ for ad-hoc retrieval.** PQ is also widely applied for ad-hoc retrieval. On one hand, the float vectors are quantized for high memory efficiency. On the other hand, the retrieval process can be significantly accelerated. For each key within database, the embedding \mathbf{z}^k is quantized as $\tilde{\mathbf{z}}_k = [\mathbf{C}_1\mathbf{b}_1, \dots, \mathbf{C}_M\mathbf{b}_M]^1$. For a given query embedding \mathbf{z}^q , the inner-product with the M codebooks can be enumerated and kept within the distance table \mathbf{T}_q , where $\mathbf{T}_q[i, j] = \langle \mathbf{z}_i^q, \mathbf{C}_{ij} \rangle$. Finally, the query and key’s inner product $\langle \mathbf{z}^q, \tilde{\mathbf{z}}^k \rangle$ can be efficiently derived by looking up the pre-computed re-

1. We adopt the Asymmetric Distance Computation (ADC) (Jégou et al., 2011), where only keys need to be quantized.

sults within \mathbf{T}_q : $\sum_{1, \dots, M} \mathbf{T}_q[i, \mathbf{b}_i]$, where no more dot product computations are needed.

2.2 Defect of Reconstruction Loss

The reconstruction loss minimization is usually adopted as the quantization model’s training objective in supervised PQ (Cao et al., 2016; Gao et al., 2019; Chen et al., 2020). It requires the distortions between key embedding and its quantization result to be reduced as much as possible:

$$\min \sum_k \|\mathbf{z}^k - \tilde{\mathbf{z}}^k\|_2. \quad (1)$$

The minimization of reconstruction loss is seemingly plausible given the following hypotheses.

Hypothesis 2.1. *The quantized keys’ embeddings are less accurate in predicting query and key’s relevance, compared with the non-quantized ones.*

Hypothesis 2.2. *The loss of prediction accuracy is a monotonously increasing function of the reconstruction loss (defined in Eq. 1).*

The first hypothesis can be “assumed correct” in reality, considering that the quantized embeddings are less expressive than the original embeddings (due to the finite number of codewords). However, the second hypothesis is problematic. In the following part, we analyze the underlying defect from both theoretical and empirical perspectives.

2.2.1 Theoretical Analysis²

We theoretically analyze two properties about the reconstruction loss: 1) it is indelible; and 2) decreasing of the reconstruction loss does not necessarily improve the prediction accuracy.

Theorem 2.1. (Positive Reconstruction Loss) *The reconstruction loss is positive if the codebooks’ scale is smaller than the key’s scale.*

That is to say, the input will always be changed after quantization given a reasonable scale of codebooks, making it impossible to keep the quantized embeddings equally expressive as the original embeddings by eliminating the reconstruction loss. Moreover, we further show that the reduction of the reconstruction loss doesn’t necessarily improve the prediction accuracy.

We start by showing the existence of “quantization invariant perturbation”, i.e., the codeword assignment will stay unchanged when such perturbations are added to the codebooks.

Lemma 2.1. *For each codebook \mathbf{C}_i , there will always exist perturbation vectors like ϵ_i , where the manipulation of codewords: $\hat{\mathbf{C}}_{i*} \leftarrow \mathbf{C}_{i*} + \epsilon_i$, doesn’t affect the codeword assignment.*

On top of the existence of quantization invariant perturbations, we may derive the “Non-monotone” about the relationship between the prediction accuracy and the reconstruction loss.

Theorem 2.2. (Non-Monotone) *PQ’s prediction accuracy is not monotonically increasing with the reduction of the reconstruction loss.*

Proof. *The statement is proved by contradiction. For each codebook, we generate a perturbation which satisfies Lemma A and add it to the codewords: $\hat{\mathbf{C}}_{i*} \leftarrow \mathbf{C}_{i*} + \epsilon_i$. According to the Lemma A, the codeword assignment will not change, so the quantized key embedding become: $\hat{\mathbf{z}}^k = \tilde{\mathbf{z}}^k + \epsilon$, where $\epsilon = [\epsilon_1, \epsilon_2, \dots, \epsilon_M]$. Now we may derive the following relationship about the reconstruction losses:*

$$\begin{aligned} \mathbb{E}\|\mathbf{z}^k - \hat{\mathbf{z}}^k\|_2 &= (\mathbb{E}(\mathbf{z}^k - \tilde{\mathbf{z}}^k)^2 + \mathbb{E}(\epsilon)^2)^{1/2} \\ &> \mathbb{E}\|\mathbf{z}^k - \tilde{\mathbf{z}}^k\|_2. \end{aligned}$$

(The first equivalence holds conditioned on the independence between $\mathbf{z}^k - \tilde{\mathbf{z}}^k$ and ϵ .)³ That is to say, the reconstruction loss is increased after the perturbation. At the same time, we may also derive the query and key’s relationship before (R) and after ($\hat{\mathbf{R}}$) the perturbation:

$$\begin{aligned} \hat{\mathbf{R}}(q, k) &= \frac{\exp(\langle \mathbf{z}^q, \hat{\mathbf{z}}^k \rangle)}{\sum_{k'} \exp(\langle \mathbf{z}^q, \hat{\mathbf{z}}^{k'} \rangle)} \\ &= \frac{\exp(\langle \mathbf{z}^q, \tilde{\mathbf{z}}^k + \epsilon \rangle)}{\sum_{k'} \exp(\langle \mathbf{z}^q, \tilde{\mathbf{z}}^{k'} + \epsilon \rangle)} \\ &= \frac{\exp(\langle \mathbf{z}^q, \tilde{\mathbf{z}}^k \rangle)}{\sum_{k'} \exp(\langle \mathbf{z}^q, \tilde{\mathbf{z}}^{k'} \rangle)} = \mathbf{R}(q, k). \end{aligned}$$

In other words, the relationship between query and key is preserved despite the increased reconstruction loss. Thus, a contradiction is obtained for the monotonous relationship between the prediction accuracy and reconstruction loss. \square

2.2.2 Empirical Analysis

To further verify the theoretical conclusions, we empirically analyze the relationship between the prediction accuracy and the reconstruction loss

2. Proofs of Theorem 2.1 and Lemma A are put to the appendix.

3. The ϵ is generated based on an arbitrary unit vector, so the independence condition always holds.

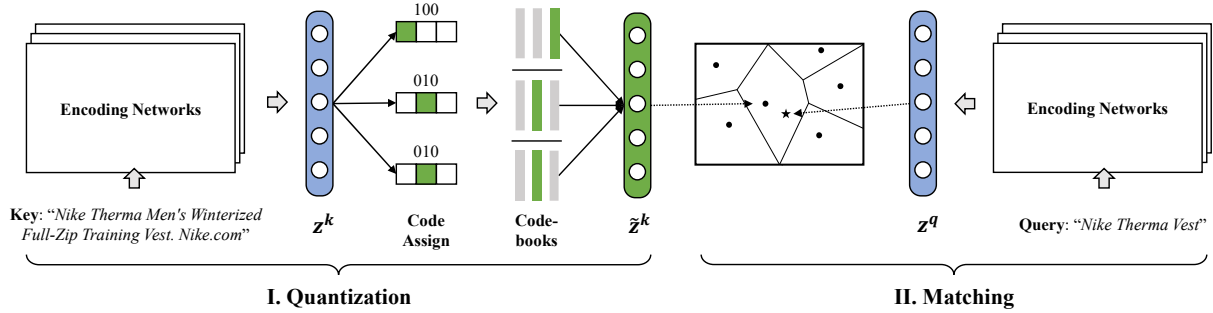


Figure 1: PQ’s retrieval workflow: (I) Quantization: the key’s embedding (z^k) is assigned to codes, whose related codewords are composed as the quantized key embedding (\tilde{z}^k); (II) Matching: the quantized key embedding (one of the centroids of the Voronoi diagram) confines the targeted query embedding z^q within its own Voronoi cell.

	Search Ads		Quora		News		Wiki	
Methods	R-loss	Recall	R-loss	Recall	R-loss	Recall	R-loss	Recall
DQN 1e-3	16.6820	0.1045	16.6647	0.5172	20.2182	0.2357	18.3501	0.0535
DQN 1e-2	7.5604	0.1433	5.6794	0.5436	9.2987	0.3138	6.9921	0.0720
DQN 1e-1	1.4021	0.1132	2.5604	0.4798	1.9051	0.2636	1.4542	0.0547
DQN 1.0	<u>0.2575</u>	0.0575	<u>0.5038</u>	0.1869	<u>0.7624</u>	0.1488	<u>0.3043</u>	0.0085

Table 1: Relationships between the reconstructive loss (R-Loss) and PQ’s accuracy (Recall@10). The reconstruction loss is monotonously reduced when its weight becomes larger. However, the smallest reconstruction loss (marked by “...”) doesn’t lead to the optimal accuracy (marked in bold).

as Table 1, by taking the deep quantization network (Cao et al., 2016) (DQN for short) as the example⁴. The weight of reconstruction loss is tuned as: 1, 1e-1, 1e-2, 1e-3 (larger weights will lead to higher loss reduction), and the weight of embedding learning is fixed to 1. We find that the reconstruction loss is monotonously reduced when a larger learning weight is used. However, the smallest reconstruction loss doesn’t bring the highest prediction accuracy (measured with Recall@10), which echos our theoretical analysis. More experimental studies in Section 4 demonstrate that the supervised PQ’s advantages over the non-supervised baselines are inconsistent and sometimes insignificant, when the reconstruction loss minimization is taken as the objective (even with the optimally tuned weights).

To summarize, the reconstruction loss cannot be eliminated with a feasible scale of codebooks, and the reduction of the reconstruction loss doesn’t necessarily improve the retrieval accuracy. As such, we turn around to formulate a new objective, where the model will stay with the reconstruction loss but optimize the PQ’s retrieval accuracy.

3 Matching-oriented PQ

We present the Matching-oriented PQ (MoPQ) in this section. In MoPQ, a new quantization objective MCL is proposed, whose minimization optimizes the query-key’s matching probability, therefore bringing about the optimal retrieval accuracy. Besides, the DCS method is introduced, which facilitates the effective minimization of MCL.

3.1 Multinoulli Contrastive Loss

The ad-hoc retrieval with PQ can be divided into two stages, shown as Figure 1. The first stage is called Quantization. The key embedding is assigned to a series of binary codes, each of which is corresponding to one codeword within a codebook. The assigned codewords are composed as the quantized key embedding, which is a centroid of the Voronoi diagram determined by the codebooks. The second stage is called Matching, where the quantized key embedding confines the targeted query embedding within its own Voronoi cell. The above matching process is probabilistically modeled as the **Multinoulli Generative Process** (as Figure 2):

- For each codebook i , a codeword is sampled from the Multinoulli distribution: $\text{Mul}(\mathbf{C}_{ij} | \mathbf{z}_i^k)$, denoted as \tilde{z}_i^k ; the quantized key embedding \tilde{z}^k is generated as the concatenation of $\{\tilde{z}_i^k\}_M$.

4. More detailed experiment settings are given in Section 4.




Figure 2: The Multinoulli Generative Process.

- The query is sampled from the distribution: $\text{Mul}(\mathbf{z}^q | \tilde{\mathbf{z}}^k)$, parameterized by the quantized key embedding $\tilde{\mathbf{z}}^k$ and query embedding \mathbf{z}^q .

Thus, the matching probability can be factorized by the joint distribution of making codeword selection from all codebooks: $\prod_i P(\mathbf{C}_{ij} | \mathbf{z}_i^k)$, and sampling query based on \mathbf{z}^q and $\tilde{\mathbf{z}}^k$: $P(\mathbf{z}^q | \tilde{\mathbf{z}}^k)$:

$$P(\mathbf{z}^q | \mathbf{z}^k) = \sum_j \prod_i P(\mathbf{C}_{ij} | \mathbf{z}_i^k) P(\mathbf{z}^q | \tilde{\mathbf{z}}^k).$$

(“ \sum_j ” indicates the enumeration of all possible codeword selection.) We expect to maximize the above query-key matching probability so as to achieve the optimal retrieval accuracy. However, the exact calculation is intractable due to: 1) the almost infinite combinations of codewords, and 2) the unknown distribution of the queries. In this place, the following transformations are made.

Firstly, we leverage the Straight Through Estimator (Bengio et al., 2013), where $P(\mathbf{C}_{ij} | \mathbf{z}_i^k)$ is transformed by the hard thresholding function:

$$P'(\mathbf{C}_{ij} | \mathbf{z}_i^k) = \begin{cases} 1, & \text{if } j = \text{argmax}\{P(\mathbf{C}_{i*} | \mathbf{z}_i^k)\}; \\ 0, & \text{otherwise.} \end{cases}$$

The above probability is calculated as $P' = (P' - P).sg() + P$ such that the gradients can be back propagated (“ $sg()$ ” is the stop gradient operation). Now the generation probability is simplified as:

$$P(\mathbf{z}^q | \mathbf{z}^k) = \prod_i P'(\mathbf{C}_{ij^*} | \mathbf{z}_i^k) P(\mathbf{z}^q | \tilde{\mathbf{z}}^k),$$

where $j^* = \text{argmax}\{P(\mathbf{C}_{i*} | \mathbf{z}_i^k)\}$. (“ \sum_j ” can be removed because $P'(\mathbf{C}_{ij} | \mathbf{z}_i^k) = 0, \forall j \neq j^*$.)

Secondly, we make a further transformation for the query’s generation probability:

$$P(\mathbf{z}^q | \tilde{\mathbf{z}}^k) = \frac{P(\tilde{\mathbf{z}}^k | \mathbf{z}^q) P(\mathbf{z}^q)}{P(\tilde{\mathbf{z}}^k)} \propto P(\tilde{\mathbf{z}}^k | \mathbf{z}^q),$$

where $P(\mathbf{z}^q)$ and $P(\tilde{\mathbf{z}}^k)$ are the prior probabilities regarded as unknown constants. The conditional

Algorithm 1: DCS Method

```

1 begin
2   for Device  $i = 1 \dots D$  do
3      $\{Q_*^i, K_*^i\}_N \leftarrow \text{Encode}(\{q_*^i, k_*^i\}_N)$ ;
4     broadcast  $\{Q_*^i, K_*^i\}_N$ ;
5   for Device  $i = 1 \dots D$  do
6      $L_p^i \leftarrow$  primary loss;
7     for Device  $j \neq i$  do
8        $L_c^{ji} \leftarrow$  image loss for Device- $j$ ;
9      $\nabla_{\theta}^i \leftarrow \nabla_{\theta}(L_p^i + \sum_j L_c^{ji})$ 
10    Update model w.r.t.  $\sum_i \nabla_{\theta}^i$ .

```

probability $P(\tilde{\mathbf{z}}^k | \mathbf{z}^q)$ calls for the normalization over the quantized key embeddings $\{\tilde{\mathbf{z}}^k\}$, which is deterministic and predefined. Now the generation probability is transformed as:

$$P(\mathbf{z}^q | \mathbf{z}^k) \propto \prod_i P'(\mathbf{C}_{ij^*} | \mathbf{z}_i^k) P(\tilde{\mathbf{z}}^k | \mathbf{z}^q) = \quad (2)$$

$$\prod_i \frac{\exp(s(\mathbf{z}_i^k, \mathbf{C}_{ij^*}))}{\sum_j \exp(s(\mathbf{z}_i^k, \mathbf{C}_{ij}))} * \frac{\exp(\langle \tilde{\mathbf{z}}^k, \mathbf{z}^q \rangle)}{\sum_{k'} \exp(\langle \tilde{\mathbf{z}}^{k'}, \mathbf{z}^q \rangle)},$$

where “ $s(\cdot)$ ” is the code assignment function (detailed forms will be discussed in Section 4.1).

The negative logarithm of the final simplification in Eq. 2 is called the **Multinoulli Contrastive Loss (MCL)**. It is used as our quantization training objective, whose minimization optimizes the query and key’s matching probability.

3.2 Approximating MCL with DCS

MCL calls for the normalization over all keys’ quantized embeddings $\{\tilde{\mathbf{z}}^k\}$, whose computation cost is huge. It has to be approximated by negative sampling (Bengio et al., 2013; Huang et al., 2020), where a subset of keys are encoded as the normalization term. Recent works (Gillick et al., 2019; Luan et al., 2020; Wu et al., 2020b; Karpukhin et al., 2020) use in-batch contrastive samples, which are free of extra encoding cost: for the i -th training instance within a mini-batch, the j -th key’s quantized embedding ($j \neq i$) will be used as a contrastive sample. Thus, there will be $N - 1$ cost free contrastive samples in total (N : batch size). Recent studies also leverage cross device in-batch sampling for sample augmentation in distributed environments (Qu et al., 2020). Particularly, a training instance on one device may take advantage of quantized key embeddings on other devices as its contrastive samples. Thus, the contrastive samples will be increased by $\times D$ times




Figure 3: Contrastive sample augmentation based on DCS. The primary loss and the image loss are combined to make the contrastive samples shared from other devices (shaded boxes) virtually differentiable.

(D : the number of devices). A problem about the cross device in-batch sampling is that the shared embeddings from other devices are not differentiable (because the shared values need to be detached from their original computation graphs). As a result, the partial gradients cannot be back-propagated through the cross-device contrastive samples, which causes distortions of the model’s update, thus undermines the optimization effect.

• **DCS Method.** We propose the **Differentiable Cross-device in-batch Sampling (DCS)**, which enables partial gradients to be back propagated for the cross-device contrastive samples. The overall workflow is summarized with Alg. 1, where the core technique is referred as the “combination of **Primary** and **Image** losses”. Suppose a total of D GPU devices are deployed, each one processes N training instances. The training instances encoded on the i -th Device are denoted as $\{Q_{1...N}^i, K_{1...N}^i\}$.⁵ The embeddings generated on each device will be broadcasted to all the other devices. As a result, the i -th device will maintain two sets of embeddings: 1) $\{Q_{1...N}^i, K_{1...N}^i\}$, which are locally encoded and differentiable, and 2) $\{Q_{1...N}^{\neq i}, K_{1...N}^{\neq i}\}$, which are broadcasted from other devices and thus non-differentiable. Each training instance (Q_j^i, K_j^i) will have two losses computed in parallel: the primary loss on the device where it is encoded (i.e., Device- i), and the image loss on the devices where it is broadcasted.

Take the first instance on Device-1 (Q_1^1, K_1^1) for illustration (as Figure 3). The query-key matching

probability $P(K_1^1|Q_1^1)$ on Device-1 is:

$$\frac{\exp(\langle Q_1^1, K_1^1 \rangle)}{\sum_j \exp(\langle Q_1^1, K_j^1 \rangle) + \sum_{\neq 1} \exp(\langle Q_1^1, \bar{K}_j^* \rangle)}. \quad (3)$$

The cross-device embeddings are detached, therefore, the partial gradients are stopped at these variables (marked as \bar{Q} and \bar{K}). The above query-key matching probability will be used by MCL (for “ $P(\bar{z}_k|z_q)$ ”) in Eq. 2, whose result is called the **primary loss** w.r.t. (Q_1^1, K_1^1) ; the sum of primary losses for all (Q_*^1, K_*^1) is denoted as L_p^1 .

Meanwhile, the query-key matching probability $P(K_1^1|Q_1^1)$ is also computed on all other devices. For the i -th device ($i \neq 1$), $P(K_1^1|Q_1^1)$ becomes:

$$\frac{\exp(\langle \bar{Q}_1^1, \bar{K}_1^1 \rangle)}{\sum_j \exp(\langle \bar{Q}_1^1, K_j^i \rangle) + \sum_{\neq i} \exp(\langle \bar{Q}_1^1, \bar{K}_j^* \rangle)}. \quad (4)$$

The differentiability is partially inverted compared with $P(K_1^1|Q_1^1)$ in Eq. 3: K_j^i becomes differentiable, but Q_1^1 and K_j^1 become non-differentiable. The above probability is used to derive another MCL, which is called the **image loss** of (Q_1^1, K_1^1) on Device- i ; the sum of image losses of all (Q_*^1, K_*^1) on Device- i is denoted as L_c^{1i} . Clearly, the above image loss will compensate the stopped gradients (related to K_j^i) in the primary loss.

The primary losses and image losses are gathered from all GPU devices and added up, based on which the model parameters θ are updated w.r.t. the following partial gradients:

$$\nabla_{\theta} \left(\sum_i (L_p^i + \sum_{j \neq i} L_c^{ij}) \right). \quad (5)$$

It can be verified that the above results are equivalent to the partial gradients derived from the fol-

5. Q_j^i and K_j^i : the query embedding (z^q) and the quantized key embedding (\bar{z}^k) of the j -th training instance on device i .

lowing full-differentiable distributions:

$$\sum_{k=1}^N \sum_{l=1}^N \frac{\exp(\langle Q_l^k, K_l^k \rangle)}{\sum_{j=1}^N \sum_{i=1}^D \exp(\langle Q_k^1, K_j^i \rangle)}. \quad (6)$$

Thus, the partial gradients are free from distortions caused by the non-differentiable variables, which enables MCL to be precisely approximated with the cross-device augmented contrastive samples.

4 Experimental Studies

4.1 Experiment Settings

- **Data.** We use three open datasets. **Quora**⁶, with question pairs of duplicated meanings (Wang et al., 2019); we use one question to retrieve its counterpart. **News**⁷, with news articles from Microsoft News (Wu et al., 2020a); we use the headline to retrieve the news body. **Wiki**⁸, with passages from Wikipedia; we use the first sentence to retrieve the remaining part of the passage. One large industrial dataset **Search Ads**, with user’s search queries and clicked ads from a worldwide search engine; we use search queries to retrieve the titles of clicked ads. (As Table 2.)

- **Baselines.** We consider both supervised and non-supervision PQ as our baselines. For supervised PQ, the embedding model and the quantization model are learned jointly. For non-supervised PQ, the embedding model is learned at first; then the quantization model is learned with fixed embeddings. We consider the following supervised methods. **DQN** (Cao et al., 2016), which is learned with two objectives: the embedding model is learned to match the query and key, and the quantization model is learned to minimize the reconstruction loss. **DVSQ** (Cao et al., 2017), which adapts the reconstruction loss to minimize the distortion of query and key’s inner product. **SPQ** (Klein and Wolf, 2019), which minimizes the disagreement between the hard and soft allocated codewords. **DPQ** (Chen et al., 2020), which still minimizes the reconstruction loss as DQN, but leverages a different quantization module. We also include 2 non-supervised baselines: the vanilla **PQ** (Jégou et al., 2011), and **OPQ** (Ge et al., 2013). Although OPQ is non-supervised, it learns the transformation of the input embeddings such that the reconstruction loss can be minimized.

	Train	Valid	Test	#Keys
News	79,122	9,891	9,891	98,388
Quora	68,240	9,055	24,041	537,340
Wiki	289,623	10,000	10,000	1,902,625
Ads	1,169,453	10,000	10,000	1,738,237

Table 2: Specifications of Datasets.

We implement two MoPQ alternatives: 1) **MoPQ^b**, the basic form with MCL and conventional in-batch sampling; 2) **MoPQ^a**, the advanced form with both MCL and DCS.

- **Implementations.** We use BERT-like Transformers (Devlin et al., 2018) for text encoding: the #layer is 4 and the hidden-dimension is 768. The input text is uncased and tokenized with WordPiece (Wu et al., 2016). The algorithms are implemented in PyTorch 1.8.0. We consider the following codeword selection functions : 1) *l*2 (default option in reality), where the codeword is selected based on Euclidean distance: $C_{ij} \leftarrow \operatorname{argmax}\{-\|\mathbf{z}_i^k - \mathbf{C}_{i*}\|_2\}$; 2) Cosine, where the codeword is selected by: $C_{ij} \leftarrow \operatorname{argmax}\{\cos(\mathbf{z}_i^k, \mathbf{C}_{i*})\}$; 3) Product, where the codeword is selected by: $C_{ij} \leftarrow \operatorname{argmax}\{\langle \mathbf{z}_i^k, \mathbf{C}_{i*} \rangle\}$; 3) Bi-linear, where the codeword is selected by: $C_{ij} \leftarrow \operatorname{argmax}\{\mathbf{z}_i^k W \mathbf{C}_{i*}^T\}$ (W is a learnable square matrix). *Our code and data will be made public-available. Pseudo codes for the algorithms, more comprehensive results and implementation details are put into supplementary materials.*

4.2 Experiment Analysis

The experimental studies focus on three major issues: 1) the overall comparisons between MoPQ and the existing PQ baselines; 2) the impact of DCS; 3) the impacts of codebook configurations, like codeword selection and codebook size.

- **Overall Comparisons.** The overall comparison results are shown in Table 3. The matching accuracy is measured by **Recall@N** (R@N). We use 8 codebooks by default, each of which has 256 codewords ($M = 8, L = 256$). The best performances are marked in bold; the most competitive baseline performances are underlined.

Firstly, it can be observed that the basic form MoPQ^b consistently outperforms all the baselines, with 7.8%~19.5% relative improvements over the most competitive baselines on different datasets. With the enhancement of DCS, MoPQ^a further improves the performances by 11.0%~42.7%, relatively. Both observations validate the effective-

6. <https://www.kaggle.com/c/quora-question-pairs>

7. <https://msnews.github.io>

8. <https://deeppgraphlearning.github.io/project/wikidata5m>

Method	Search Ads			Quora			News			Wiki		
	R10	R50	R100	R10	R50	R100	R10	R50	R100	R10	R50	R100
MoPQ ^b	0.2141	0.3513	0.4193	0.5861	0.8381	0.9054	0.3422	0.5462	0.6334	0.1005	0.2150	0.2878
MoPQ ^a	0.2439	0.3820	0.4563	0.6717	0.8976	0.9416	0.3799	0.5787	0.6533	0.1434	0.2658	0.3322
PQ	0.1252	0.2126	0.2682	0.3503	0.6353	0.7518	0.2028	0.4040	0.5069	0.0340	0.1015	0.1496
OPQ	0.1506	0.2507	0.3127	0.4572	0.7349	0.8359	0.2381	0.4277	0.5180	0.0558	0.1409	0.2082
DQN	0.1433	0.2520	0.3145	<u>0.5436</u>	<u>0.8285</u>	<u>0.8977</u>	<u>0.3138</u>	<u>0.5217</u>	<u>0.6108</u>	0.0720	0.1728	0.2451
DVSQ	0.1548	0.2784	0.3498	0.4061	0.7097	0.8179	0.3045	0.5105	0.6102	<u>0.0841</u>	<u>0.1849</u>	<u>0.2602</u>
SPQ	<u>0.1907</u>	<u>0.3145</u>	<u>0.3844</u>	0.5321	0.7984	0.8782	0.2470	0.4474	0.5449	0.0655	0.1648	0.2306
DPQ	0.1709	0.2953	0.3650	0.4941	0.7760	0.8621	0.2552	0.4641	0.5639	0.0599	0.1561	0.2254

Table 3: Comparisons between MoPQ (upper), non-supervised (middle) and supervised (lower) PQ baselines.

Method	Ads	Quora	News	Wiki
PQ	21.4227	13.6768	16.1511	16.9847
DQN	7.5604	5.6794	9.2987	6.9921
MoPQ ^b	35.7987	29.7244	26.4847	26.5861
MoPQ ^a	39.5576	30.0319	27.0710	26.8888

Table 4: Comparison of reconstruction loss.

ness of our proposed methods. As for the baselines: the supervised PQ’s performances are comparatively higher than the non-supervised ones; however, the improvements are not consistent: in some cases, OPQ achieves comparable or even higher recall rates than some of the supervised PQ.

Secondly, we further clarify the relationship between the reconstruction loss and the query-key matching accuracy. More results on reconstruction loss are reported in Table 4. For one thing, we find that by jointly learning the embedding and quantization models, DQN’s reconstruction losses become significantly lower than PQ. At the same time, DQN’s recall rates are consistently higher than PQ. Such observations indicate that: to some extent, the reduction of reconstruction loss may help to improve PQ’s retrieval accuracy. On the other hand, the reconstruction losses of MoPQ^b and MoPQ^a are much higher than DQN, but it dominates the baselines in terms of recall rate. Such observations echo the theoretical and empirical findings in Section 2.2: PQ’s query-key retrieval accuracy will not monotonously increase with the reduction of reconstruction loss.

A brief conclusion for the above observations: although the minimization of reconstruction loss still contributes to PQ’s retrieval accuracy, the improvement is limited due to the non-monotone between both factors. In contrast, the minimization of MCL directly maximizes the query-key’s matching probability, which makes MoPQ achieve much more competitive retrieval accuracy.

- **Impact of DCS.** More analysis about DCS

is shown in the upper of Table 5. We consider two baselines: 1) the conventional in-batch sampling, where no cross-device sampling is made; 2) the Non-differentiable Cross-device Sampling (NCS), which also makes embeddings shared across GPU devices for contrastive sample augmentation, but no image loss is computed to compensate the stopped gradients. It is observed that both DCS and NCS outperform the in-batch sampling, thanks to the augmentation of contrastive samples. However, the improvement of NCS is limited compared with DCS. As discussed in Section 3.2, the gradients are stopped at the non-differentiable variables of NCS, which causes distortions for the model’s update. Thus, the model’s training outcome is restricted because of it.

- **Impact of codeword selection.** We use MoPQ^b as the representative to analyze the impact of different forms of codeword selections in the lower of Table 5. We find that MoPQ’s performances are not sensitive to the codeword selection function, as the experiment results are quite close to each other. Given that the l_2 selection’s performance is slightly better and no extra parameters are introduced, it is used as our default choice.

- **Impact of codebook size.** We analyze the impact of codebook size in Table 6; the SPQ baseline is included for comparison. With the expansion of scale, i.e., more codebooks and more codewords per codebook, MoPQ’s performance can be improved gradually. In all of the settings, MoPQ maintains its advantage over SPQ. It is also observed that in some cases, MoPQ outperforms SPQ with even smaller codebook size, e.g., MoPQ (M=4, L=256) and SPQ (M=8, L=256) on News; in other words, the higher recall rate is achieved with smaller space and time costs.

Method	Search Ads			Quora			News			Wiki		
	R10	R50	R100	R10	R50	R100	R10	R50	R100	R10	R50	R100
In-Batch	0.2141	0.3513	0.4193	0.5861	0.8381	0.9054	0.3422	0.5462	0.6334	0.1005	0.2150	0.2878
NCS	0.2249	0.3620	0.4305	0.6264	0.8655	0.9197	0.3569	0.5534	0.6355	0.1151	0.2320	0.2960
DCS	0.2439	0.3820	0.4563	0.6717	0.8976	0.9416	0.3799	0.5787	0.6533	0.1434	0.2658	0.3322
l_2	0.2141	0.3513	0.4193	0.5861	0.8381	0.9054	0.3422	0.5462	0.6334	0.1005	0.2150	0.2878
Cosine	0.2019	0.3297	0.3987	0.5804	0.8296	0.8993	0.3215	0.5160	0.6031	0.0784	0.1828	0.2518
Product	0.2027	0.3290	0.3998	0.5578	0.8001	0.8707	0.3256	0.5307	0.6158	0.0843	0.1930	0.2617
Bi-Lin	0.2102	0.3410	0.4113	0.5878	0.8382	0.9034	0.3197	0.5210	0.6116	0.1115	0.2259	0.2965

Table 5: The impact of DCS (upper table), and the impact of codeword selection functions (lower table).

Method	Ads	Quora	News	Wiki
SPQ (M=4, L=128)	0.1018	0.2804	0.1569	0.0461
MoPQ ^b (M=4, L=128)	0.1068	0.4484	0.2327	0.0563
SPQ (M=4, L=256)	0.1264	0.3006	0.1833	0.0440
MoPQ ^b (M=4, L=256)	0.1289	0.4648	0.2649	0.0569
SPQ (M=8, L=256)	0.1907	0.5321	0.2470	0.0655
MoPQ ^b (M=8, L=256)	0.2141	0.5861	0.3422	0.1005
SPQ (M=16, L=256)	0.2196	0.5767	0.2780	0.1029
MoPQ ^b (M=16, L=256)	0.2574	0.6395	0.3889	0.1457

Table 6: Impact of codebook size. (Recall@10)

5 Conclusion

In this paper, we propose MoPQ to optimize PQ’s ad-hoc retrieval accuracy. A systematic revisit is made for the existing supervised PQ, where we identify the limitation of using reconstruction loss minimization as the training objective. We propose MCL as our new training objective, where the model can be learned to maximize the query-key matching probability to achieve the optimal retrieval accuracy. We further leverage DCS for contrastive sample argumentation, which ensures the effective minimization of MCL. The experiment results on 4 real-world datasets validate the effectiveness of our proposed methods.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (No. U1936104) and The Fundamental Research Funds for the Central Universities 2020RC25.

References

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.

Yue Cao, Mingsheng Long, Jianmin Wang, and Shichen Liu. 2017. Deep visual-semantic quantization for efficient image retrieval. In *Proceedings of*

the IEEE Conference on Computer Vision and Pattern Recognition, pages 1328–1337.

Yue Cao, Mingsheng Long, Jianmin Wang, Han Zhu, and Qingfu Wen. 2016. Deep quantization network for efficient image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, page 3457–3463.

Ting Chen, Lala Li, and Yizhou Sun. 2020. Differentiable product quantization for end-to-end embedding compression. In *International Conference on Machine Learning*, pages 1617–1626.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Lianli Gao, Xiaosu Zhu, Jingkuan Song, Zhou Zhao, and Heng Tao Shen. 2019. Beyond product quantization: Deep progressive quantization for image retrieval. *arXiv preprint arXiv:1906.06698*.

Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. 2013. Optimized product quantization. *IEEE transactions on pattern analysis and machine intelligence*, 36(4):744–755.

Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldridge, Eugene Ie, and Diego Garcia-Olano. 2019. Learning dense representations for entity retrieval. In *Proceedings of the 23rd Conference on Computational Natural Language Learning*, pages 528–537.

Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. 2020. Embedding-based retrieval in facebook search. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 2553–2561.

Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128.

Hervé Jégou, Romain Tavenard, Matthijs Douze, and Laurent Amsaleg. 2011. Searching in one billion vectors: re-rank with source coding. In *2011 IEEE*

International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 861–864.

Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6769–6781.

Benjamin Klein and Lior Wolf. 2019. End-to-end supervised product quantization for image search and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5041–5050.

Yiding Liu, Guan Huang, Jiayang Liu, Weixue Lu, Suqi Cheng, Yukun Li, Daiting Shi, Shuaiqiang Wang, Zhicong Cheng, and Dawei Yin. 2021. Pre-trained language model for web-scale retrieval in baidu search. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 3365–3375.

Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2020. Sparse, dense, and attentional representations for text retrieval. *arXiv preprint arXiv:2005.00181*.

Jingwei Ma, Jiahui Wen, Mingyang Zhong, Weitong Chen, and Xue Li. 2019. Mmm: Multi-source multi-net micro-video recommendation with clustered hidden item representation learning. *Data Science and Engineering*, pages 240–253.

Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2020. Rocketqa: An optimized training approach to dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2010.08191*.

Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2019. Kepler: A unified model for knowledge embedding and pre-trained language representation. *arXiv preprint arXiv:1911.06136*.

Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, and Ming Zhou. 2020a. MIND: A large-scale dataset for news recommendation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3597–3606.

Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020b. Scalable zero-shot entity linking with dense entity retrieval. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6397–6407.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus

Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

Mingdong Zhu, Derong Shen, Lixin Xu, and Xianfang Wang. 2021. Scalable multi-grained cross-modal similarity query with interpretability. *Data Science and Engineering*, pages 280–293.

A Proof

Theorem 2.1. (Positive Reconstruction Loss)

The reconstruction loss is positive if the codebooks’ scale is smaller than the key’s scale.

Proof. The statement is proved by induction. First of all, it’s trivial that the reconstruction loss is positive when the number of codewords is one ($L = 1$). Suppose we are given codebooks \mathbf{C} , where $L > 1$ and the corresponding reconstruction loss is positive. Now assign a new codeword $\mathbf{C}_{*,L+1}$ to each codebook, where $[\mathbf{C}_{1,L+1}, \dots, \mathbf{C}_{M,L+1}] = \mathbf{z}^k$ (\mathbf{z}^k is an arbitrary key’s embedding). With the augmentation of codebooks, the reconstruction loss related to k will become zero, which means the reconstruction loss will be reduced by $\|\mathbf{z}^k - \tilde{\mathbf{z}}^k\|_2$. This is to say, when the size of codebooks is smaller than the number of keys, the reconstruction loss can always be reduced by increasing the codebooks’ scale. Considering that the reconstruction loss is lower bounded by 0, it is obvious that the reconstruction loss is always positive given finite codebooks. \square

Lemma 2.1. For each codebook \mathbf{C}_i , there will always exist perturbation vectors like ϵ_i , where the manipulation of codewords: $\hat{\mathbf{C}}_{i*} \leftarrow \mathbf{C}_{i*} + \epsilon_i$, doesn’t affect the codeword assignment.

Proof. The existence of ϵ_i is proved by giving one “always-valid” example. We define the following thresholding radius for the perturbation:

$$r_{\epsilon_i} \leftarrow 0.5 * \min\{\min\{\|\mathbf{C}_{il} - \mathbf{z}_i\|_2 : l \neq j\} - \|\mathbf{C}_{ij} - \mathbf{z}_i\|_2 : \forall \mathbf{z}\},$$

where \mathbf{C}_{ij} is the assigned codeword for i -th sub-vector \mathbf{z}_i of embedding \mathbf{z} , i.e., $j = \operatorname{argmin}\{\|\mathbf{C}_{i*} - \mathbf{z}_i\|_2\}$.

Let $r \sim \operatorname{Uniform}(0, r_{\epsilon_i})$, and $\epsilon_i \leftarrow r * \mathbf{u}_i$ (\mathbf{u}_i is an arbitrary unit vector). In this way, $\forall l \neq j$:

$$2\|\epsilon_i\|_2 < 2r_{\epsilon_i} \leq \|\mathbf{C}_{il} - \mathbf{z}_i\|_2 - \|\mathbf{C}_{ij} - \mathbf{z}_i\|_2 \quad (7)$$

Add this ϵ_i to all codewords in the codebook \mathbf{C}_i : $\hat{\mathbf{C}}_{i*} \leftarrow \mathbf{C}_{i*} + \epsilon_i$. According to the inequation (7)

M&N	Method	Search Ads			Quora			News			Wiki		
		R10	R50	R100	R10	R50	R100	R10	R50	R100	R10	R50	R100
M=4 N=128	PQ	0.0248	0.0633	0.0895	0.1255	0.3214	0.4368	0.1455	0.3084	0.4002	0.0302	0.0872	0.1402
	OPQ	0.0365	0.0890	0.1252	0.2134	0.3214	0.4368	0.1388	0.3215	0.4218	0.0213	0.0676	0.1048
	DQN	0.0319	0.0846	0.1227	0.3618	0.6722	0.7759	0.0895	0.2250	0.3119	0.0241	0.0647	0.0975
	DVSQ	0.0390	0.1005	0.1483	0.3855	0.7097	0.8118	0.1628	0.3460	0.4529	0.0413	0.0703	0.1383
	SPQ	0.1018	0.2166	<u>0.2882</u>	0.2804	0.6095	0.7373	0.1569	0.3378	0.4424	0.0440	0.1098	0.1621
	DPQ	0.0454	0.1371	0.1980	0.2956	0.6431	0.7729	0.1434	0.3223	0.4260	0.0458	0.0927	0.1345
	MoPQ ^b	<u>0.1068</u>	<u>0.2199</u>	0.2876	<u>0.4484</u>	<u>0.7472</u>	<u>0.8417</u>	<u>0.2327</u>	<u>0.4445</u>	<u>0.5419</u>	<u>0.0563</u>	<u>0.1428</u>	<u>0.2040</u>
	MoPQ ^a	0.1342	0.2567	0.3266	0.5130	0.8012	0.8815	0.2496	0.4447	0.5372	0.0742	0.1702	0.2341
M=4 N=256	PQ	0.0348	0.0802	0.1388	0.1673	0.3893	0.5136	0.1455	0.3084	0.4002	0.0337	0.1071	0.1475
	OPQ	0.0533	0.1144	0.1581	0.2556	0.5253	0.6549	0.1640	0.3561	0.4634	0.0277	0.0851	0.1307
	DQN	0.0391	0.1040	0.1485	0.4421	0.7430	0.8341	0.1267	0.2892	0.3868	0.0320	0.0846	0.1238
	DVSQ	0.0481	0.1287	0.1884	0.4103	0.7446	0.8446	0.2011	0.4080	0.5163	0.0447	0.1269	0.1842
	SPQ	0.1264	0.2416	0.3066	0.3006	0.6142	0.7360	0.1833	0.3754	0.4679	0.0461	0.1130	0.1635
	DPQ	0.0630	0.1627	0.2272	0.3111	0.6513	0.7688	0.1666	0.3606	0.4652	0.0493	0.0986	0.1446
	MoPQ ^b	<u>0.1289</u>	<u>0.2627</u>	<u>0.3109</u>	<u>0.4648</u>	<u>0.7603</u>	<u>0.8545</u>	<u>0.2649</u>	<u>0.4721</u>	<u>0.5674</u>	<u>0.0569</u>	<u>0.1415</u>	<u>0.2009</u>
	MoPQ ^a	0.1506	0.2802	0.3486	0.5285	0.8057	0.8817	0.2904	0.4909	0.5778	0.0854	0.1902	0.2615
M=16 N=256	PQ	0.2184	0.3413	0.4123	0.4520	0.7870	0.8715	0.3093	0.5269	0.6194	0.0769	0.1993	0.2718
	OPQ	0.2260	0.3588	0.4227	0.3788	0.5854	0.6693	0.2381	0.4277	0.5180	0.1160	0.2322	0.3073
	DQN	0.2347	0.3718	0.4414	0.6061	0.8538	0.9100	0.3667	0.5629	0.6478	0.1411	0.2784	0.3533
	DVSQ	0.2257	0.3489	0.4156	0.5819	0.8352	0.9046	0.3327	0.5420	0.6271	0.1291	0.2538	0.3245
	SPQ	0.2196	0.3574	0.4258	0.5767	0.8354	0.8943	0.2780	0.4867	0.5798	0.1029	0.2210	0.2964
	DPQ	0.2196	0.3541	0.4222	0.5254	0.7990	0.8781	0.3241	0.5260	0.6105	0.0952	0.2102	0.2845
	MoPQ ^b	<u>0.2574</u>	<u>0.3893</u>	<u>0.4608</u>	<u>0.6395</u>	<u>0.8732</u>	<u>0.9291</u>	<u>0.3889</u>	<u>0.5814</u>	<u>0.6596</u>	<u>0.1654</u>	<u>0.2946</u>	<u>0.3697</u>
	MoPQ ^a	0.2722	0.4176	0.4874	0.7075	0.9118	0.9506	0.4411	0.6217	0.6909	0.2065	0.3411	0.4077

Table 7: Performance for different codebooks’ scale. The best performances are marked in bold, and the second are underlined.

and triangle inequality, we can get:

$$\begin{aligned}
\|\hat{\mathbf{C}}_{ij} - \mathbf{z}_i\|_2 &\leq \|\mathbf{C}_{ij} - \mathbf{z}_i\|_2 + \|\epsilon_i\|_2 \\
&< \|\mathbf{C}_{il} - \mathbf{z}_i\|_2 - \|\epsilon_i\|_2 \\
&< \|\hat{\mathbf{C}}_{il} - \mathbf{z}_i\|_2,
\end{aligned}$$

where $l \neq j$. Therefore, it’s obvious that:

$$j = \operatorname{argmin}\{\|\hat{\mathbf{C}}_{i*} - \mathbf{z}_i\|_2\}.$$

The \mathbf{z}_i will still be mapped to j -th codeword in codebook \mathbf{C}_i . In other words, the original codeword assignment will not be affected by ϵ_i . \square

B Training Details

The models are implemented with PyTorch 1.8.0 and run with $8 \times$ Nvidia-A100-40G GPUs. We use an average pooling over the last transformers layer as the text embedding and then apply an additional linear layer to reduce the size of embedding to 128. The weight of reconstruction loss in baselines is tuned from $\{1, 1e-1, 1e-2, 1e-3, 1e-4\}$. Without explicit mention, we default use eight codebooks, each of which has 256 codewords. We optimize the parameters with the Adam optimizer. The learning rate is $1e-4$ for pretrained transformers layers and $1e-3$ for other layers (e.g., product quantization). We train the models for 50 epochs with a batch size of 500. We evaluate the model

every epoch on the validation set and keep the best checkpoint for the final evaluation on test sets.

C More Experiments Results

We conduct extensive experiments for MoPQ and baselines on four datasets. The results are shown in Table 7. The best performances with the same settings are marked in bold, and the second are underlined. The observations from this table are consistent with the conclusions in our paper:

- The performance of PQ models can be improved by increasing the number of codebooks and codewords.
- MoPQ consistently obtains better performance than baselines.
- Adopting DCS method to augment contrastive samples, MoPQ^a achieves significant improvement than MoPQ^b.