

# TL;DR: TWIN LEARNING FOR DIMENSIONALITY REDUCTION

Anonymous authors

Paper under double-blind review




Figure 1: **Overview of the proposed TLDR, a dimensionality reduction method.** Given a set of feature vectors in a generic input space, we use nearest neighbors to define a set of feature pairs whose proximity we want to preserve. We then learn a dimensionality-reduction function (the *encoder*) by encouraging neighbors in the input space to have similar representations. We learn it jointly with an auxiliary *projector* that produces high dimensional representations, where we compute the Barlow Twins (Zbontar et al., 2021) loss over the  $d' \times d'$  cross-correlation matrix averaged over the batch.

## ABSTRACT

Dimensionality reduction methods are unsupervised approaches which learn low-dimensional spaces where some properties of the initial space, typically the notion of “neighborhood”, are preserved. They are a crucial component of diverse tasks like visualization, compression, indexing, and retrieval. Aiming for a totally different goal, self-supervised visual representation learning has been shown to produce transferable representation functions by learning models that encode invariance to artificially created distortions, *e.g.* a set of hand-crafted image transformations. Unlike manifold learning methods that usually require propagation on large  $k$ -NN graphs or complicated optimization solvers, self-supervised learning approaches rely on simpler and more scalable frameworks for learning.

In this paper, we unify these two families of approaches from the angle of manifold learning and propose TLDR, a dimensionality reduction method for generic input spaces that is porting the simple self-supervised learning framework of Zbontar et al. (2021) to a setting where it is hard or impossible to define an appropriate set of distortions by hand. We propose to use nearest neighbors to build pairs from a training set and a redundancy reduction loss borrowed from the self-supervised literature to learn an encoder that produces representations invariant across such pairs. TLDR is a method that is simple, easy to implement and train, and of broad applicability; it consists of an offline nearest neighbor computation step that can be highly approximated, and a straightforward learning process that does not require mining negative samples to contrast, eigendecompositions, or cumbersome optimization solvers. Aiming for scalability, the Achilles’ heel of manifold learning, we focus on improving *linear* dimensionality reduction, a technique that is still an integral part of many large-scale systems. By simply replacing PCA with TLDR, we are able to increase the performance of GeM-AP (Revaud et al., 2019), a state-of-the-art landmark recognition method by 4% mAP for 128 dimensions, and to retain its performance with  $16\times$  fewer dimensions.

## 1 INTRODUCTION

Dimensionality reduction refers to a set of unsupervised approaches which aim at learning low-dimensional spaces where properties of an initial higher-dimensional input space, *e.g.* proximity or “neighborhood”, are preserved. It is a crucial component for very diverse tasks, ranging from visualization and compression, to indexing and retrieval; most web-scale retrieval systems, still use dimensionality reduction in practice. Assuming that data in the input space lie on a lower-dimensional “manifold”, dimensionality reduction is also referred to as manifold learning.

Recently, and aiming for a different goal, self-supervised representation learning has been shown to produce representations that are highly transferable to a wide number of downstream tasks via encoding invariance to image distortions like data augmentations (Chen et al., 2020a; He et al., 2020; Caron et al., 2020; Zbontar et al., 2021; Grill et al., 2020). Central to the success of such methods is the scalable and easy-to-optimize learning framework that such methods adopt, often based on loss functions with or without contrasting pairs. Seeing how manifold learning methods lack scalability and usually require propagation on large  $k$ -NN graphs or complicated optimization solvers, one cannot help but wonder: *Can we borrow from the highly successful learning frameworks of self-supervised representation learning to design dimensionality reduction approaches?*

In this paper, we unify these two families of approaches from the angle of manifold learning and propose *Twin Learning for Dimensionality Reduction* or **TLDR**, a generic dimensionality-reduction technique where the only prior is that data lies on a reliable manifold we want to preserve. It is based on the intuition that comparing a data point and its nearest neighbors is a good “distortion” to learn from, and hence a good way of approximating the local manifold geometry. Similar to other manifold learning methods (Roweis & Saul, 2000; Van der Maaten & Hinton, 2008; Belkin & Niyogi, 2003; Donoho & Grimes, 2003; Hadsell et al., 2006) we use Euclidean nearest neighbors as a way of defining distortions of the input that the dimensionality reduction function should be invariant to. However, unlike other manifold learning methods, TLDR does not require eigendecompositions, negatives to contrast, or cumbersome optimization solvers; it simply consists of an offline nearest neighbor computation step that can be highly approximated without loss in performance and a straightforward stochastic gradient descent learning process. This leads to a highly scalable method that can learn linear and non-linear encoders for dimensionality reduction while trivially handling out-of-sample generalization. We show an overview of the proposed method in Figure 1.

We are interested in explicitly targeting applications like image and document search where training labels are non-existent and dimensionality reduction is an important part of the state-of-the-art pipelines. Aiming at large-scale search applications, we focus on improving *linear* dimensionality reduction with a compact encoder, an integral part of the first-stage of most retrieval systems, and an area where PCA (Pearson, 1901) is still the default method used in practice (Revaud et al., 2019; Tolia et al., 2020). We present a large set of ablations and experimental results on two common benchmarks for image retrieval (Radenović et al., 2018a), as well as on the natural language processing task of argument retrieval. We show that one can achieve significant gains without altering the encoding and search complexity: for example we can improve landmark image retrieval on ROxford5K (Radenović et al., 2018a) by almost 4 *mAP* points for 128 dimensions, a commonly used dimensionality (Tolia et al., 2020), when replacing PCA with TLDR in a state-of-the-art method (Revaud et al., 2019).

**Contributions.** We introduce TLDR, a dimensionality reduction method that achieves neighborhood embedding learning with the simplicity and effectiveness of recent self-supervised visual representation learning losses. Aiming for scalability, we focus on large-scale image and document retrieval where dimensionality reduction is still an integral component. We show that replacing PCA (Pearson, 1901) with a *linear* TLDR encoder can greatly improve the performance of state-of-the-art methods without any additional computational complexity. We thoroughly ablate parameters and show that our design choices allow TLDR to be robust to a large range of hyper-parameters and is applicable to a diverse set of tasks and input spaces. *We intend to make the code for TLDR publicly available.*

## 2 TWIN LEARNING FOR DIMENSIONALITY REDUCTION

Starting from a set of unlabeled and high-dimensional features, our goal is to learn a lower-dimensional space which preserves the local geometry of the larger input space. Assuming that we have no prior knowledge other than the reliability of the local geometry of the input space, we use

nearest neighbors to define a set of feature pairs whose proximity we want to preserve. We then learn the parameters of a dimensionality-reduction function (the *encoder*) using a loss that encourages neighbors in the input space to have similar representations, while also minimizing the redundancy between the components of these vectors. Similar to other works (Chen et al., 2020b;a; Zbontar et al., 2021) we append a *projector* to the encoder that produces a representation in a very high dimensional space, where the Barlow Twins (Zbontar et al., 2021) loss is computed. At the end of the learning process, the projector is discarded. All aforementioned components are detailed next. We call our method **Twin Learning for Dimensionality Reduction** or **TLDR**, in homage to the Barlow Twins loss. An overview is provided in Figure 1.

**Preserving local neighborhoods.** Recent self-supervised learning methods define positive pairs via hand-crafted distortions that exploit prior information from the input space. In absence of any such prior knowledge, defining local distortions can only be achieved via assumptions on the input manifold. Assuming a locally linear manifold, for example, would allow using the Euclidean distance as a local measure of on-manifold distortion and using nearest neighbors over the training set would be a good approximation for local neighborhoods. Therefore, we construct pairs of neighboring training vectors, and learn invariance to the distortion from one such vector to another. Practically, we define the local neighborhood of each training sample as its  $k$  nearest neighbors. Although defining local neighborhood in such a uniform way over the whole manifold might seem naive, we experimentally show that not only it is sufficient, but also that our algorithm is robust across a wide range of  $k$  values (see Section 3.1).

Using nearest neighbors is of course not the only way of defining neighborhoods. In fact, we show alternative results with a simplified variant of TLDR (denoted as TLDR<sub>G</sub>) where we construct pairs by simply adding Gaussian noise to an input vector. This is a baseline resembling denoising autoencoders, although in our case we are using a) an asymmetric encoder-decoder architecture and b) the Barlow twins loss instead of a reconstruction loss.

**Notation.** Our goal is to learn an encoder  $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^d$  that takes as input a vector  $x \in \mathbb{R}^D$  and outputs a corresponding reduced vector  $z = f_\theta(x) \in \mathbb{R}^d$ , with  $d \ll D$ . Without loss of generality, we define the encoder to be a neural network with learnable parameters  $\theta$ . Let  $\mathcal{X}$  be a (training) set of datapoints in  $\mathbb{R}^D$ , the  $D$ -dimensional input space. Let  $x \in \mathbb{R}^D$  be a vector from  $\mathcal{X}$ .  $\mathcal{N}_k(x)$  is composed of the  $k$  nearest neighbors of  $x$ . For a vector  $y \in \mathcal{X}$  from the training set:  $y \in \mathcal{N}_k(x) \Leftrightarrow y \in \arg_k \min_{y \in \mathcal{X}} d(x, y)$ , where  $d(\cdot, \cdot)$  denotes the Euclidean distance. Although the definition above can be trivially extended to non-Euclidean distances and adaptive neighborhoods (e.g. defined by a radius), without loss of generality we present our method and results with pairs from  $k$  Euclidean neighbors. We define *neighbor pairs* as pairs  $(x, y) \in \mathcal{X} \times \mathcal{X}$  where  $y \in \mathcal{N}_k(x)$ .

**Learning à la Barlow Twins.** Although contrastive losses were proven highly successful for visual representation learning, explicitly minimizing the redundancy of the output dimension is highly desirable for dimensionality reduction: having a highly informative output space is more important than a highly discriminative one. We therefore choose to learn the parameters of our encoder by minimizing the Barlow Twins loss function (Zbontar et al., 2021), that suits perfectly. Similar to (Zbontar et al., 2021), we append a projector  $g_\phi$  to the encoder  $f_\theta$ , allowing to calculate the loss in a (third) representation space which is not the one that will be used for subsequent tasks. That extended space can possibly be much larger. We detail the encoder and the projector later. Let  $\hat{z} = g_\phi(f_\theta(x))$  be the output vector of the projector,  $\hat{z} \in \mathbb{R}^{d'}$ . Given a pair of neighbors  $(x^A, x^B)$  and the corresponding vectors  $\hat{z}^A, \hat{z}^B$  after the projector, the loss function  $\mathcal{L}_{BT}$  is given by:

$$\mathcal{L}_{BT} = \sum_i (1 - C_{ii})^2 + \lambda \sum_i \sum_{i \neq j} C_{ij}^2, \quad \text{where } C_{ij} = \frac{\sum_b \hat{z}_{b,i}^A \hat{z}_{b,j}^B}{\sqrt{\sum_b (\hat{z}_{b,i}^A)^2} \sqrt{\sum_b (\hat{z}_{b,j}^B)^2}}, \quad (1)$$

where  $b$  indexes the positive pairs in a batch,  $i$  and  $j$  are two dimensions from  $\mathbb{R}^{d'}$  (i.e.  $0 \leq i, j \leq d'$ ) and  $\lambda$  is a hyper-parameter.  $C$  is the  $d' \times d'$  cross-correlation matrix computed and averaged over all positive pairs  $(\hat{z}^A, \hat{z}^B)$  from the current batch. The loss is composed of two terms. The first term encourages the diagonal elements to be equal to 1. This makes the learned representations invariant to applied distortions, i.e. the datapoints moving along the input manifold in the neighborhood of a training vector are encouraged to share similar representations in the output space. The second term is pushing off-diagonal elements towards 0, reducing the redundancy between output dimensions, a highly desirable property for dimensionality reduction.

The redundancy reduction term can be viewed as a soft-whitening constraint on the representations and, shown in Zbontar et al. (2021), it works better than performing “hard” whitening on the representations (Ermolov et al., 2021). Finally, it is worth noting that understanding the dynamics of learning without contrasting pairs is far from trivial and beyond the scope of this paper; we refer the reader to the recent work by Tian et al. (2021) that studies this learning paradigm in depth and discusses why trivial solutions are avoided when learning without negatives as in Eq. (1).

**The encoder  $f_\theta$ .** We consider a number of different architectures for the encoder:

- *linear*: The most straight-forward choice for encoder  $f_\theta$  is a linear function parametrized by a  $D \times d$  weight matrix  $W$  and bias term  $b$ , i.e.  $f_\theta(x) = Wx + b$ . Beyond computational benefits, and given that we are mostly interested in medium-sized output spaces where  $d \in \{8, \dots, 512\}$ , we argue that, given a meaningful enough input space, a linear encoder could suffice in preserving neighborhoods of the input.
- *factorized linear*: Exploiting the fact that batch normalization (BN) (Ioffe & Szegedy, 2015) is linear during inference,<sup>1</sup> we formulate  $f_\theta$  as a multi-layer linear model, where  $f_\theta$  is a sequence of  $l$  layers, each composed of a linear layer followed by a BN layer. This model introduces non-linear dynamics which can potentially help during training but the sequence of layers can still be replaced with a single linear layer after training for efficiently encoding new features.
- *MLP*:  $f_\theta$  can be a multi-layer perceptron with batch normalization (BN) (Ioffe & Szegedy, 2015) and rectified linear units (reLUs) as non-linearities, i.e.  $f_\theta$  would be a sequence of  $l$  linear-BN-reLU triplets, each with  $H^i$  hidden units ( $i = 1, \dots, l$ ), followed by a linear projection.

Our main goal is to develop a scalable alternative to PCA for dimensionality reduction, so we are mostly interested in linear and factorized linear encoders. It is worth already mentioning that, as we will show in our experimental validation, gains from introducing an MLP in the encoder are minimal and would not justify the added computational cost in practice.

**The projector  $g_\phi$ .** As also recently noted in Tian et al. (2021), a crucial part of learning with non-contrastive pairs is the projector. This module is present in a number of contrastive self-supervised learning methods (Chen et al., 2020a; Grill et al., 2020; Zbontar et al., 2021; Tian et al., 2021). It is usually implemented as an MLP inserted between the transferable representations and the loss function. Unlike other methods, however, where the projector takes the representations to an even lower dimensional space for the contrastive loss to operate on (i.e. for SimCLR (Chen et al., 2020a) and BYOL (Grill et al., 2020),  $d' \ll d$ ), for the Barlow Twins objective, operating in large output dimensions is crucial.

In Section 3, we study the impact of the dimension  $d'$  and experiment with a wide range of values. We empirically verify the findings of Zbontar et al. (2021) that calculating the de-correlation loss in higher dimensions ( $d' \gg d$ ) is highly beneficial. In this case, and as shown in Figure 1, the transferable representation is now the *bottleneck* layer of this non-symmetrical hour-glass model. Although Eq. (1) is applied after the projector and only indirectly decorrelates the output representation components, having more dimensions to decorrelate leads to a representation that is more informative: the *bottleneck* effect created by the projector’s output being in a much larger dimensionality implicitly enables the network to learn an encoder that also has more decorrelated outputs.

### 3 EXPERIMENTAL VALIDATION

In this section, we present a set of experiments validating the proposed TLDR both on visual and textual features. We selected one input representation space and one task for each modality: for the visual domain, we focus on the task of *landmark image retrieval* (Section 3.1) and use 2048-dimensional global image features from an off-the-shelf ResNet50 pre-trained for image retrieval<sup>2</sup> (Revaud et al., 2019). For the textual domain, we focus on the task of *argument retrieval* (Section 3.2). We use 768-dimensional features from an off-the-shelf Bert-Siamese model called

<sup>1</sup>Although batch normalization is non-linear during training because of the reliance on the current batch statistics, during inference and using the means and variances accumulated over training, it reduces to a linear scaling applied to the features, that can be embedded in the weights of an adjacent linear layer.

<sup>2</sup><https://github.com/naver/deep-image-retrieval>

ANCE<sup>3</sup> (Xiong et al., 2021) trained for document retrieval, following the dataset definitions from Thakur et al. (2021). Details on tasks and datasets used are summarized in Table A in the Appendix.

**Implementation details.** We do not explicitly normalize representations during learning TLDR; yet, we follow the common protocol and L2-normalize the features before retrieval for both tasks. Results reported for PCA use whitened PCA; we tested multiple whitening power values and kept the ones that performed best. Further implementation details are reported in the Appendix. It is noteworthy that we used the *exact same hyper-parameters* for the learning rate, weight decay, scaling, and  $\lambda$  suggested in Zbontar et al. (2021), despite having very different tasks and encoder architectures.

**Further ablations, results on FashionMNIST and 2D visualizations.** Due to space constraints, additional results and interesting ablations can be found in the Appendix. In particular, we explore the effect of the training set size, the batch size and report results on another NLP task: duplicate query retrieval. Moreover, and although beyond the scope of what TLDR is designed for, in Appendix E we present results on FashionMNIST when using TLDR on raw pixel data and for 2D visualization. We show that for cases where the input pixels are forming an informative space, TLDR can achieve top performance for  $d \geq 8$ .

### 3.1 RESULTS ON LANDMARK IMAGE RETRIEVAL

We first focus on landmark image retrieval. For large-scale experiments on this task, it is common practice to apply dimensionality reduction to global normalized image representations using PCA with whitening (Jégou & Chum, 2012; Tolias et al., 2016; Revaud et al., 2019). We start from GeM-AP (Revaud et al., 2019) and simply replace the standard PCA step with our proposed TLDR.

**Experimental protocol.** We start from 2048-dimensional features obtained from the pre-trained ResNet-50 of (Revaud et al., 2019), which uses Generalized-Mean pooling (Radenović et al., 2018b) and has been specifically trained for landmark retrieval using the AP loss (GeM-AP). To learn the dimensionality reduction function, we use a dataset composed of 1.5 million landmark images (Weyand et al., 2020). We learn different output spaces whose dimensions range from 32 to 512. Finally, we evaluate these spaces on two standard image retrieval benchmarks (Radenović et al., 2018a), the revisited Oxford and Paris datasets (ROxford5K and RParis6K). Each dataset comes with two test sets of increasing difficulty, the “Medium” and “Hard”. Following these datasets’ protocol, we apply the learned dimensionality reduction function to encode both the gallery images and the set of query images whose 2048-dim features have been extracted beforehand with the model of Revaud et al. (2019). We then evaluate landmark image retrieval on ROxford5K and RParis6K and report mean average precision (mAP), the standard metric reported for these datasets. For brevity, we report the “Mean” mAP metric, *i.e.* the average of the mAP of the “Medium” and “Hard” test sets; we include the individual plots for “Medium” and “Hard” in the Appendix for completeness.

**Compared approaches.** We report results for several flavors of our approach. *TLDR* uses a linear projector, *TLDR*<sub>1</sub> uses a factorized linear one, and *TLDR*<sub>1</sub><sup>\*</sup> an MLP encoder with 1 hidden layer. As an alternative, we also report *TLDR*<sub>G</sub>, which uses Gaussian noise to create synthetic neighbour pairs. All variants use an MLP with 2 hidden layers and 8192 dimensions as a projector.

We compare with a number of un- and self-supervised methods (see also Table B in the Appendix for a summary). First, and foremost, we compare to reducing the dimension with *PCA* with whitening, which is still standard practice for these datasets (Revaud et al., 2019; Radenović et al., 2018b; Tolias et al., 2020). We also report results for our approach but trained with the Mean Square Error reconstruction loss instead of the Barlow Twins’ (as we discuss in Section 4, *PCA* can be rewritten as learning a linear encoder and projector with a reconstruction loss), and refer to this method as *MSE*. In this case, the projector’s output is reduced to 2048 dimensions in order to match the input’s dimensionality. Following a number of approaches that use nearest neighbors as (self-)supervision for contrastive learning (Hadsell et al., 2006), the *Contrastive* approach uses a contrastive loss on top of the projector’s output. This draws inspiration from Hadsell et al. (2006), and is a variant where we replace the Barlow Twins loss, with the loss from Hadsell et al. (2006). It is worth noting that we omit results from a more faithful reimplementation of Hadsell et al. (2006), *i.e.* using a max-margin loss directly applied on the lower dimensional space and without a projector, as they were very low. Note that none of the manifold learning method we tested was able to neither scale, nor outperform *PCA*

<sup>3</sup>[https://www.sbert.net/docs/pretrained\\_models.html](https://www.sbert.net/docs/pretrained_models.html)

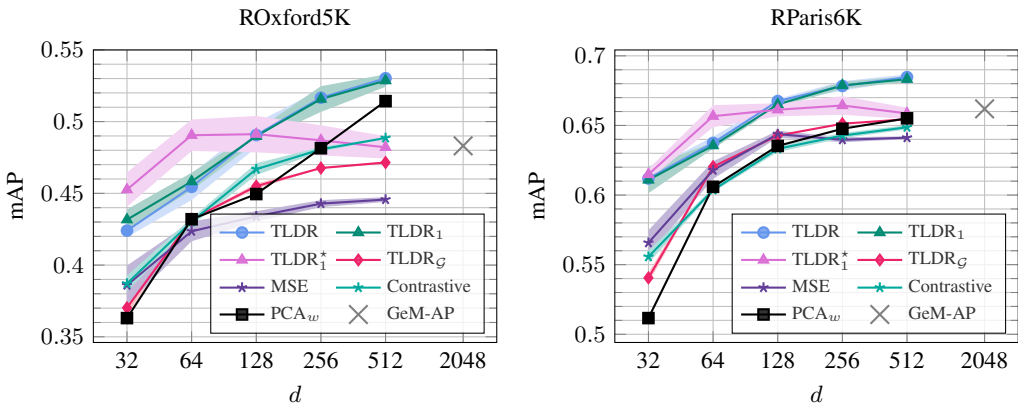


Figure 2: **Image retrieval experiments.** Mean average precision (mAP) on ROxford5K (left) and RParis6K (right) as a function of the output dimensions  $d$ . We report TLDR with different encoders: linear (TLDR), factorized linear with 1 hidden layer (TLDR<sub>1</sub>), and a MLP with 1 hidden layer (TLDR<sub>1</sub><sup>\*</sup>), the projector remains the same (MLP with 2 hidden layers). We compare with PCA with whitening, two baselines based on TLDR, but which respectively train with a reconstruction (MSE) and a contrastive (Contrastive) loss, and also with  $TLDR_G$ , a variant of TLDR which uses Gaussian noise to synthesize pairs. The original GeM-AP performance is also reported.

in output dimensions  $d \geq 8$ ; we present comparisons for smaller  $d$  in Section 3.3. Finally, we report retrieval results obtained on the initial features from [Revaud et al. \(2019\)](#) (*GeM-AP*), *i.e.* without dimensionality reduction. For all flavours of TLDR, we fix the number of nearest neighbors to  $k = 3$ , although, and as we show in Figure 4, TLDR performs well for a wide range of number of neighbors.

**Results.** Figure 2 reports mean average precision (mAP) results for ROxford5K and RParis6K; as the output dimensions  $d$  varies. We report the average of the Medium and Hard protocols for brevity, while results per protocol are presented in Appendix C.1. We make a number of observations. First and most importantly, we observe that both linear flavors of our approach outperform PCA by a significant margin. For instance, TLDR improves ROxford5K retrieval by almost 4 mAP points for 128 dimensions over the PCA baseline. The MLP flavor is very competitive for very small dimensions (up to 128) but degrades for larger ones. Even for the former, it is not worth the extra-computational cost. An important observation is that we are able to retain the performance of the input representation (GeM-AP) while using only 1/16th of its dimensionality. Using a different loss (MSE and Contrastive) instead of the Barlow Twins’ in TLDR degrades the results. These approaches are comparable to or worse than PCA. Finally, replacing true neighbors by synthetic ones, as in  $TLDR_G$ , performs worse.

### 3.2 RESULTS ON FIRST STAGE DOCUMENT RETRIEVAL

For document retrieval, the process is generally divided into two stages: the first one selects a small set of candidates while the second one re-ranks them. Because it works on a smaller set, this second stage can afford costly strategies, but the first stage has to scale. The typical way to do this is to reduce the dimension of the representations used in the first retrieval stage, often in a supervised fashion ([Khattab & Zaharia, 2020](#); [Gao et al., 2021](#)). Following our initial motivation, we investigate the use of unsupervised dimensionality reduction for document retrieval scenarios where a supervised approach is not possible, *e.g.* when no such training data is available.

**Experimental protocol.** We start from 768-dimensional features extracted from a model trained for Question Answering (QA), *i.e.* ANCE ([Xiong et al., 2021](#)). We use Webis-Touché-2020 ([Bondarenko et al., 2020](#); [Wachsmuth et al., 2017](#)) a conversational argument dataset composed of 380k documents to learn the dimensionality reduction function.

**Compared approaches.** We report results for three flavors of our approach. *TLDR* uses a linear encoder while *TLDR<sub>1</sub>* and *TLDR<sub>2</sub>* use a factorized linear one with respectively one hidden layer and two hidden layers. We compare with PCA, which was the best performing competitor from Section 3.1. We also report retrieval results obtained with the 768-dimensional initial features.

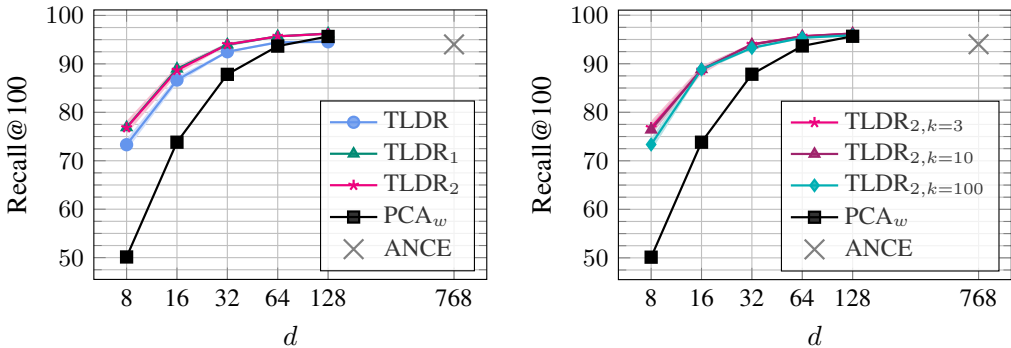


Figure 3: **Argument retrieval results on ArguAna** for different values of output dimensions  $d$ . On the left we vary the amount of factorized layers, with fixed  $k = 3$ , on the right we fix the amount of factorized layers to 2 and test  $k = [3, 10, 100]$ . Factorized linear is fixed to 512 hidden dimensions.

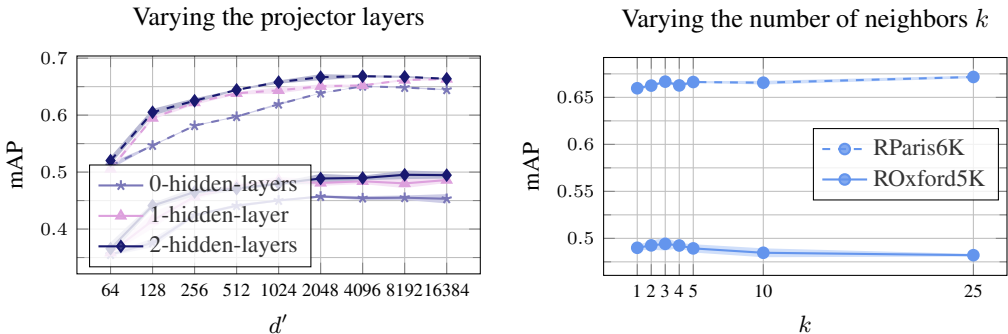


Figure 4: **Impact of TLDR hyper-parameters** with a linear encoder and  $d = 128$ . Dashed (solid) lines are for RParis6K-Mean (ROxford5K-Mean). (Left) Impact of the *auxiliary* dimension  $d'$  and the number of hidden layers in the projector. (Right) Impact of the number of neighbors  $k$ . We see how the algorithm is robust to the number of neighbors used.

**Results.** Figure 3 reports retrieval results on ArguAna, for different output dimensions  $d$ . We observe that the linear version of TLDR outperforms PCA for almost all values of  $d$ . The linear-factorized ones outperforms PCA in all scenarios. We see that the gain brought by TLDR over PCA increases as  $d$  decreases. Note that we achieve results equivalent to the initial ANCE representation using only 4% of the original dimensions; PCA, needs twice as many dimensions to achieve similar performance.

### 3.3 ANALYSIS AND IMPACT OF HYPER-PARAMETERS

**Impact of hyper-parameters.** Figure 4 studies the role of some of our parameters. On the left side of Figure 4, we vary the architecture of the projector  $g_\phi$ , an important module of TLDR. We see that having hidden layers generally helps. As also noted in Zbontar et al. (2021), having a high auxiliary dimension  $d'$  for computing the loss is very important and highly impacts performance. On the right side of Figure 4 we show the surprisingly consistent performance of TLDR across a wide range of numbers of neighbors  $k$ . We observe the same stability across several batch sizes (see also Figure E).

**Comparisons to manifold learning methods on smaller output dimensions.** In Figure 5a we present results for TLDR when the output dimensionality is  $d' \leq 64$ ; in this regime, a few more manifold learning methods can be run, eg UMAP, Locally Linear Embedding (LLE) (Roweis & Saul, 2000), Local Tangent Space Alignment (LTSA) (Zhang & Zha, 2004), and UMAP (McInnes et al., 2018). Unfortunately, even at smaller output dimensions we had to subsample the dataset to run some of the methods, due to their scalability issues. Specifically, we are forced to use only 5% of the training set ( $\sim 75K$  images) for learning LLE and LTSA, and 50% ( $\sim 750K$  images) for UMAP.

**How sensitive is TLDR to approximate nearest neighbors?** To verify that our system is robust to an approximate computation of nearest neighbors, we test its performance using product quan-

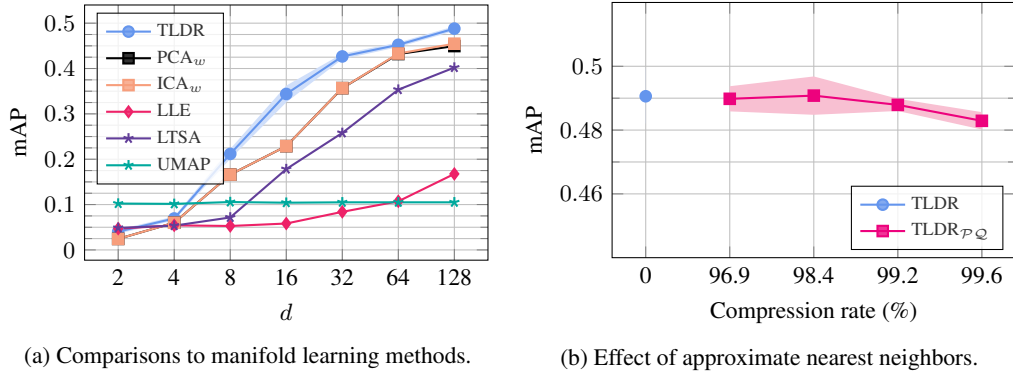


Figure 5: Left: Comparisons to manifold learning methods for small output dimensions  $d \leq 128$ . Mean average precision (mAP) on ROxford5K (Radenović et al., 2018a) averaged over the Medium and Hard test sets as a function of the output dimensions  $d$ . Right: The effect of nearest neighbor approximation for  $d = 128$ . We plot mAP as a function of the embedding compression rate used during nearest neighbor computation. Note that the baseline (compression rate = 0) is using the 2048-dimensional (8192 bytes) GeM-AP representations during nearest neighbor computation.

tization (Ge et al., 2013) while varying the quantization budget (i.e. the amount of bytes used for each image during the nearest neighbor search). Compression is done using optimized product quantization (OPQ) (Ge et al., 2013) via the FAISS library (Johnson et al., 2017) and results are reported in Figure 5b. We see that TLDR is quite robust to quantization during the nearest neighbor search and that even when the quantization is pretty strong (1/64 the default size or merely 16 Bytes per vector) TLDR still retains its high performance.

#### 4 DISCUSSION AND RELATED WORK

The basic idea behind TLDR is embarrassingly simple and links to a large number of related methods, from PCA to manifold learning and neighborhood embedding. In this section we discuss a few such relations; more are discussed in Appendix F due to lack of space.

**Linear dimensionality reduction.** We refer the reader to Cunningham & Ghahramani (2015) for an extensive review of linear dimensionality reduction. It is beyond the scope of this paper to exhaustively discuss many such related works, we will therefore focus on PCA (Pearson, 1901) which is the de facto standard linear dimensionality method, in particular for large-scale retrieval.

One can derive the learning objective of PCA (Pearson, 1901) by setting  $f_\theta(x) = W^T x$  and  $g(x) = Wx$  in the model of Figure 1, i.e. use a linear encoder and projector with  $W \in \mathbb{R}^{D \times d}$ , and optimize  $W$  via minimizing the Frobenius norm of the matrix of reconstruction errors over the whole training set, subject to orthogonality constraints:

$$W^* = \arg \min_W \|x - g(f_\theta(x))\|_F, \quad \text{s.t.} \quad W^T W = I_d. \quad (2)$$

This equation has a closed form solution that can be obtained via the eigendecomposition of the data covariance matrix and then keeping the largest  $d$  eigenvectors.

Unlike PCA, TLDR does not constrain the projector to be a linear model, nor the loss to be a reconstruction loss. In fact, the redundancy reduction term in the Barlow Twins loss encourages the *whitening* of the batch representations as a soft constraint (Zbontar et al., 2021), in a way analogous to the orthogonality constraint of Eq.(2). We see from Figure 4 that part of the performance gains of TLDR over PCA is precisely due to this asymmetry in the architecture, i.e. when the projector is an MLP with hidden layers. Looking at MSE results in Figures 2, i.e. a version of TLDR with a reconstruction loss, we also see that the Barlow Twins loss and the flexibility of computing it in an arbitrarily high  $d'$ -dimensional space further contributes to this gain. One can therefore interpret TLDR as a more generic way of optimizing a linear encoder, i.e. using an arbitrary decoder and approximating the constraint of Eq.(2) in a soft way, further incorporating a weak notion of whitening.



**Manifold learning and neighborhood embedding methods.** Manifold learning methods define objective functions that try to preserve the local structure of the input manifold, usually expressed via a  $k$ -NN graph. Non-linear unsupervised dimensionality-reduction methods usually require the  $k$ -NN graph of the input data, while most further require eigenvalue decompositions (Roweis & Saul, 2000; Donoho & Grimes, 2003; Zhang & Zha, 2004) and shortest-path (Tenenbaum et al., 2000) or computation of the graph Laplacian (Belkin & Niyogi, 2003). Others involve more complex optimization (McInnes et al., 2018; Agrawal et al., 2021). Moreover, many manifold learning methods were created to solely operate on the data they were learned on. Although “out-of-sample” extensions for many of such methods have been proposed (Bengio et al., 2004), methods like Spectral Embeddings, pyMDE (Agrawal et al., 2021) or the very popular  $t$ -SNE (Van der Maaten & Hinton, 2008) can only be used for the data they were trained on. Finally, UMAP (McInnes et al., 2018) was recently proposed as not only a competitor of  $t$ -SNE on 2-dimensional outputs, but as a general purpose dimension reduction technique. Yet, all our experiments with UMAP, even after exhaustive hyperparameter tuning, resulted in very low performance for  $d \geq 8$  for all the tasks we evaluated in the main paper and the Appendix.

**Nearest neighbors as “supervision” for contrastive learning.** The seminal method DrLIM (Hadsell et al., 2006) uses a contrastive loss over neighborhood pairs for representation learning. Experimenting only on simple datasets like MNIST, it learns a CNN backbone and the dimensionality reduction function in a single stage, using a max-margin loss. TLDR resembles DrLIM (Hadsell et al., 2006) with respect to the encoder input and the way pairs are constructed; a crucial difference, however, is the loss function and the space in which it is computed: DrLIM uses a contrastive loss which is computed directly on the lower dimensional space. Despite our best effort to make this approach work as described, performance was very low without a projector. Using the contrastive loss from Hadsell et al. (2006), together with the projector we use for TLDR, we were able to get more meaningful results (reported as Contrastive in our experiments), although still underperforming the Barlow Twins loss. This difference may be due to two reasons: first, and as discussed above, Barlow Twins encourages the whitening of the representations which makes it more suitable for this task. Second, and like many other pair-wise losses, the contrastive loss further requires sampling hard/meaningful negatives (Wu et al., 2017; Radenović et al., 2018b).

**TLDR out of its comfort zone.** Although TLDR can be seen as a way of generalizing recent self-supervised visual representation learning methods to cases where handcrafted transformations of the data are challenging or impossible to define, we want to emphasize that it *is not suited* for self-supervised representation learning from pixels; augmentation invariance is a much more suited prior in that regard, while it is also practically impossible to define meaningful neighboring pairs from the input pixel space. Additionally, although visualization is a common manifold learning application, TLDR is neither designed nor recommended for 2D outputs; there are other methods like Van der Maaten & Hinton (2008); McInnes et al. (2018); Agrawal et al. (2021) that specialize for such tasks.

**What is TLDR suitable for?** TLDR excels for dimensionality reduction to mid-size outputs, *e.g.* when  $d$  is from 32 to 256 dimensions. This is very useful in practice for retrieval and a set of output dimensions where the vast majority of manifold learning methods cannot scale. At the same time, TLDR enables the community to utilize a powerful learning framework initially tailored for visual representation learning (Zbontar et al., 2021) in different domains like natural language.

## 5 CONCLUSIONS

In this paper we introduce TLDR, a dimensionality-reduction method that combines neighborhood embedding learning with the simplicity and effectiveness of recent self-supervised learning losses. By simply replacing PCA with TLDR one can significantly increase the state-of-the-art landmark retrieval performance of GeM-AP (Revaud et al., 2019) and boost argument retrieval performance without additional computational cost. TLDR further offers a number of desirable properties: i) *Scalability*: learned via stochastic gradient descent, TLDR can easily be parallelized across GPUs and machines, while for even the largest datasets, approximate nearest neighbor methods can be used to create input pairs in sub-linear complexity (Ge et al., 2013; Kalantidis & Avrithis, 2014), ii) *Simplicity*: The Barlow Twins (Zbontar et al., 2021) objective is robust and easy to optimize, and does not have trivial solutions, iii) *Out-of-sample generalization*, and iv) *Linear encoding complexity*: TLDR is highly effective with a linear encoder, offering a direct replacement of PCA without extra encoding cost.

## REPRODUCIBILITY STATEMENT

We report all the hyperparameters we used and all implementation details needed to reproduce our experiments in Section 3, and Appendices C and D. We report the urls for the publicly available pre-trained models we used to extract the input features. All the datasets we use are publicly available to download. Finally, we also intend to make easy-to-use code for TLDR publicly available.

## REFERENCES

- Akshay Agrawal, Alnur Ali, and Stephen Boyd. Minimum-distortion embedding. *arXiv preprint arXiv:2103.02559*, 2021.
- Ehsan Amid and Manfred K Warmuth. Trimap: Large-scale dimensionality reduction using triplets. *arXiv preprint arXiv:1910.00204*, 2019.
- Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural codes for image retrieval. In *Proc. ECCV*, 2014.
- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- Yoshua Bengio, Jean-François Paiement, Pascal Vincent, Olivier Delalleau, Nicolas Le Roux, and Marie Ouimet. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. *Proc. NeurIPS*, 16: 177–184, 2004.
- Alexander Bondarenko, Maik Fröbe, Meriem Beloucif, Lukas Gienapp, Yamen Ajjour, Alexander Panchenko, Chris Biemann, Benno Stein, Henning Wachsmuth, Martin Potthast, et al. Overview of touché 2020: Argument retrieval. In *Proc. CLEF*, pp. 384–395. Springer, 2020.
- Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *Proc. ICML*, pp. 1597–1607, 2020a.
- Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020b.
- John P Cunningham and Zoubin Ghahramani. Linear dimensionality reduction: Survey, insights, and generalizations. *JMLR*, 16(1):2859–2900, 2015.
- David L Donoho and Carrie Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proc. PNAS*, 100(10):5591–5596, 2003.
- Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In *Proc. ICML*. PMLR, 2021.
- Zhiyuan Fang, Jianfeng Wang, Lijuan Wang, Lei Zhang, Yezhou Yang, and Zicheng Liu. {SEED}: Self-supervised distillation for visual representation. In *Proc. ICLR*, 2021.
- Luyu Gao, Zhuyun Dai, and Jamie Callan. Coil: Revisit exact lexical match in information retrieval with contextualized inverted list. *arXiv preprint arXiv:2104.07186*, to appear in *NAACL21*, 2021.
- Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. Optimized product quantization. *PAMI*, 36(4):744–755, 2013.
- Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. Deep image retrieval: Learning global representations for image search. In *Proc. ECCV*, pp. 241–257. Springer, 2016.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. In *Proc. NeurIPS*, 2020.
- Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proc. KDD*, pp. 855–864, 2016.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *Proc. CVPR*, volume 2, 2006.

- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proc. CVPR*, pp. 9729–9738, 2020.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Doris Hoogeveen, Karin Verspoor, and Timothy Baldwin. Cqadupstack: Gold or silver. In *Proc. SIGIR*, volume 16, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proc. ICML*, 2015.
- Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, Teddy Furon, and Ondřej Chum. Efficient diffusion on region manifolds: Recovering small objects with compact cnn representations. In *Proc. CVPR*, pp. 2077–2086, 2017.
- Ahmet Iscen, Yannis Avrithis, Giorgos Tolias, Teddy Furon, and Ondřej Chum. Fast spectral ranking for similarity search. In *Proc. CVPR*, pp. 7632–7641, 2018a.
- Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Mining on manifolds: Metric learning without labels. In *Proc. CVPR*, pp. 7642–7651, 2018b.
- Hervé Jégou and Ondřej Chum. Negative evidences and co-occurrences in image retrieval: the benefit of PCA and whitening. In *Proc. ECCV*, pp. 774–787. Springer, 2012.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017.
- Yannis Kalantidis and Yannis Avrithis. Locally optimized product quantization for approximate nearest neighbor search. In *Proc. CVPR*, pp. 2321–2328, 2014.
- Omar Khattab and Matei Zaharia. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In *Proc. SIGIR*, 2020.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. Distilling dense representations for ranking using tightly-coupled teachers. *arXiv preprint arXiv:2010.11386*, 2020.
- Chundi Liu, Guangwei Yu, Maksims Volkovs, Cheng Chang, Himanshu Rai, Junwei Ma, and Satya Krishna Gorti. Guided similarity separation for image retrieval. In *Proc. NeurIPS*, 2019.
- Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human generated machine reading comprehension dataset. In *CoCo@ NeurIPS*, 2016a.
- Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. Ms marco: A human generated machine reading comprehension dataset. In *CoCo@ NeurIPS*, 2016b.
- Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proc. CVPR*, pp. 3967–3976, 2019.
- Karl Pearson. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proc. KDD*, pp. 701–710, 2014.
- Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting Oxford and Paris: Large-scale image retrieval benchmarking. In *Proc. CVPR*, pp. 5706–5715, 2018a.
- Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning CNN image retrieval with no human annotation. *PAMI*, 41(7):1655–1668, 2018b.
- J. Revaud, J. Almazan, R.S. Rezende, and C.R. de Souza. Learning with average precision: Training image retrieval with a listwise loss. In *Proc. ICCV*, 2019.

- Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Learning local feature descriptors using convex optimisation. *PAMI*, 36(8):1573–1585, 2014.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th international conference on world wide web*, pp. 1067–1077, 2015.
- Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. BEIR: a heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663*, 2021.
- Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019.
- Yuandong Tian, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning dynamics without contrastive pairs. In *Proc. ICML*, 2021.
- Giorgos Tolias, Ronan Sifre, and Hervé Jégou. Particular object retrieval with integral max-pooling of CNN activations. In *Proc. ICLR*, 2016.
- Giorgos Tolias, Tomas Jeníček, and Ondřej Chum. Learning and aggregating deep local descriptors for instance-level recognition. In *Proc. ECCV*, pp. 460–477. Springer, 2020.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *JMLR*, 9(11), 2008.
- Henning Wachsmuth, Martin Potthast, Khalid Al Khatib, Yamen Ajjour, Jana Puschmann, Jiani Qu, Jonas Dorsch, Viorel Morari, Janek Bevendorff, and Benno Stein. Building an argument search engine for the web. In *Proceedings of the 4th Workshop on Argument Mining*, pp. 49–59, 2017.
- Henning Wachsmuth, Shahbaz Syed, and Benno Stein. Retrieval of the best counterargument without prior topic knowledge. In *Proc. ACL*, pp. 241–251, 2018.
- T. Weyand, A. Araujo, B. Cao, and J. Sim. Google Landmarks Dataset v2 - A Large-Scale Benchmark for Instance-Level Recognition and Retrieval. In *Proc. CVPR*, 2020.
- C-Y Wu, Manmatha R., Smola A.J., and Krahenbuhl P. Sampling matters in deep embedding learning. In *Proc. ICCV*, 2017.
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *Proc. ICLR*, 2021.
- Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. When does self-supervision help graph convolutional networks? In *Proc. ICML*, pp. 10871–10880, 2020.
- Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow Twins: Self-supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021.
- Zhenyue Zhang and Hongyuan Zha. Principal manifolds and nonlinear dimensionality reduction via tangent space alignment. *SIAM journal on scientific computing*, 26(1):313–338, 2004.

# Appendix

## Table of Contents

---

<b>A Appendix Summary</b>	<b>13</b>
<b>B Tables of tasks and compared approaches</b>	<b>14</b>
<b>C Additional experiments: Landmark image retrieval</b>	<b>14</b>
C.1 Results on Medium and Hard protocols separately . . . . .	14
C.2 Results with “Oracle” nearest neighbors . . . . .	14
C.3 ResNet-101 features . . . . .	14
C.4 Varying the size of the training set . . . . .	14
C.5 Batch size ablation . . . . .	15
<b>D Additional experiments: Document retrieval</b>	<b>15</b>
D.1 Tasks and dataset . . . . .	16
D.2 Experimental results . . . . .	18
<b>E FashionMNIST: Learning from raw pixel data and visualization</b>	<b>19</b>
<b>F Further discussions and related works</b>	<b>20</b>
F.1 Limitations of our work . . . . .	22

---

## A APPENDIX SUMMARY

In this appendix we present a number of additional details, results and Figures that we could not fit in the main text due to lack of space. In summary:

- We present tables with the **tasks and datasets** we explore in the main paper (Table A), as well as with a **summary of all compared methods** (Table B).
- We report additional experiments for the **landmark retrieval** task in Appendix C. Specifically, we present results for the med/hard splits separately (Appendix C.1), an experiment with oracle neighbors (Appendix C.2), an experiment with features from a larger ResNet-101 backbone (Appendix C.3).
- We present ablations when **varying the training set size** (Appendix C.4) and the **batch size** (Appendix C.5).
- We report additional experiments for the **document retrieval** in Appendix D. Specifically, we extend our evaluation protocol and report result on a new task: **duplicate query retrieval**. We further investigate not only dimensionality reduction for the same task, but also the case of **dimensionality reduction transfer**.
- Although TLDR is not suited for such applications, as a proof of concept we present results on the FashionMNIST dataset in Appendix E, *i.e.* when **learning from raw pixel data**. We also present some results when using TLDR for **visualization**, *i.e.* when the output dimension is  $d = 2$  in Figure K.
- We extend Section 4 with **further discussion** on related topics and more related works in Appendix F. We conclude with a brief discussion on limitations of TLDR in Appendix F.1

Task (Metric)	Input feature space	Dimensionality reduction dataset	Test dataset
Landmark Retrieval (mAP)	ResNet50 features $D = 2048$ trained on Landmarks-clean (40k) (Babenko et al., 2014; Gordo et al., 2016)	Google Landmarks (Weyand et al., 2020) (1.5M)	ROxford5K (5k) RParis6K (6k) (Radenović et al., 2018a)
Argument Retrieval (Recall@100)	BERT Features $D = 768$ trained on MSMarco (8.8M) (Nguyen et al., 2016a)	Webis-Touché 2020 (380k) (Bondarenko et al., 2020)	ArguAna (3k) (Wachsmuth et al., 2018)

Table A: **Datasets and tasks** of the main paper.

Method	(Self-) supervision	Encoder	Projector	Loss	Notes
PCA (Pearson, 1901)	unsupervised	linear	linear	Reconstruction MSE + orthogonality	Used for dimensionality reduction in SoTA methods like DELF, GeM, GeM-AP and HOW
DrLim Contrastive	neighbor-supervised neighbor-supervised	MLP linear	None MLP	Contrastive Contrastive	(Hadsell et al., 2006) ( <i>very low performance</i> ) Hadsell et al. (2006) with projector
MSE TLDR <sub>G</sub>	unsupervised denoising	linear linear	MLP MLP	Reconstruction MSE Barlow Twins	TLDR with MSE loss TLDR with noise as distortion
TLDR TLDR <sub>1,2</sub> TLDR <sub>1,2</sub> <sup>+</sup>	neighbor-supervised neighbor-supervised neighbor-supervised	linear fact. linear MLP	MLP MLP MLP	Barlow Twins Barlow Twins Barlow Twins	

Table B: **Compared Methods.** For *unsupervised* methods the objective is based on reconstruction, *neighbor-supervised* methods utilize nearest neighbors as pseudo-labels to learn, *denoising* learns to ignore added Gaussian noise.

## B TABLES OF TASKS AND COMPARED APPROACHES

### C ADDITIONAL EXPERIMENTS: LANDMARK IMAGE RETRIEVAL

#### C.1 RESULTS ON MEDIUM AND HARD PROTOCOLS SEPARATELY

In Figure A we report the mAP metric for the Medium and Hard splits of the Revisited Oxford and Paris datasets (Radenović et al., 2018a) separately.

#### C.2 RESULTS WITH “ORACLE” NEAREST NEIGHBORS

In Figure B we present results using an oracle version of TLDR, *i.e.* a version that uses labels to only keep as pairs neighbors that come from the same landmark in the training set. As we see, TLDR practically matches the oracle’s performance.

#### C.3 RESNET-101 FEATURES

We also experimented with features obtained from a larger pre-trained ResNet-101 model from Revaud et al. (2019)<sup>4</sup>. We see in Figure C that TLDR retains a significant gain over PCA and in fact surpasses the highest state-of-the-art numbers based on global features as reported in Toliás et al. (2020) for ROxford5K.

#### C.4 VARYING THE SIZE OF THE TRAINING SET

In Figure D we show the impact of the size of the training set on TLDR’s performance by randomly selecting subsets of images of increasing size from the Google Landmarks training set (Weyand et al., 2020). As we see, PCA outperforms TLDR for a reduced number of images, however, it does not benefit from adding more data, keeping the same performance across all training set sizes. In contrast, TLDR does benefit from adding more data; all plots suggest that a larger training set could potentially boost the performance even further, increasing the gap with respect to PCA.

<sup>4</sup><https://github.com/naver/deep-image-retrieval>

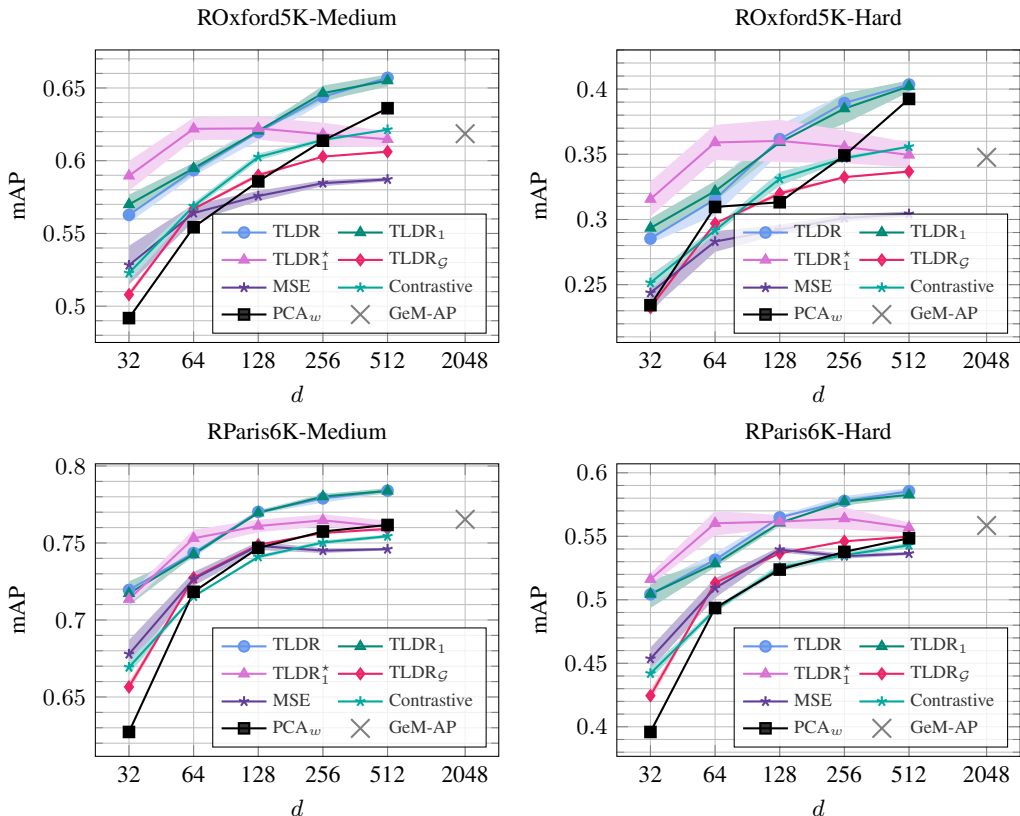


Figure A: **Image retrieval experiments.** Mean average precision (mAP) on ROxford5K (top) and RParis6K (bottom), for the Medium (left) and Hard (right) test sets, as a function of the output dimensions  $d$ . We report TLDR with different encoders: linear (TLDR), factorized linear with 1 hidden layer (TLDR<sub>1</sub>), and a MLP with 1 hidden layer (TLDR<sub>1</sub><sup>+</sup>), the projector remains the same (MLP with 2 hidden layers). We compare with two baselines based on TLDR, but which respectively train with a reconstruction (MSE) and a contrastive (Contrastive) loss. Our main baselines are PCA with whitening, and the original 2048-dimensional features (GeM-AP [Revaud et al. \(2019\)](#)), *i.e.* before projection.

### C.5 BATCH SIZE ABLATION

Finally, in Figure E, we show results of TLDR varying the size of the training mini-batch. Surprisingly, we observe it is stable across a wide range of values, allowing training TLDR under limited memory resources.

## D ADDITIONAL EXPERIMENTS: DOCUMENT RETRIEVAL

In Section 3 of the main paper we studied first stage document retrieval under the task of argument retrieval, where both dimensionality reduction and evaluation are performed on datasets designed for the same task. In this section, we extend this evaluation protocol introducing a new task: duplicate query retrieval and now investigate not only dimensionality reduction for the same task, but also the case of *dimensionality reduction transfer*.

In the following paragraphs we first introduce the five datasets we use for first stage document retrieval, and then we discuss the additional experiments involving duplicate query datasets.




Figure B: **Neighbor-supervised with oracle.** Mean average precision (mAP) on ROxford5K (Radenović et al., 2018a) for the Medium (left) and Hard (right) test sets, as a function of the output dimensions  $d$ . We compare TLDR with an oracle version that uses labels to select training pairs. We include as baselines both PCA and ICA with whitening, and the original 2048-dimensional features (GeM-AP [32]), i.e. before projection.

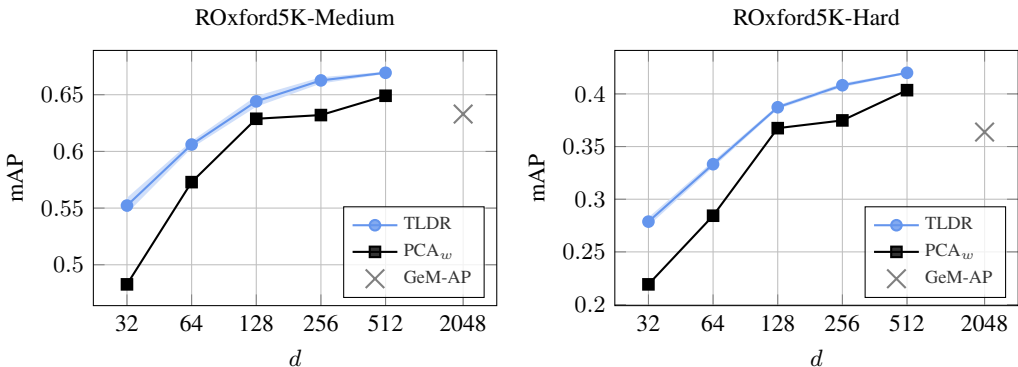


Figure C: **ResNet-101 features.** Mean average precision (mAP) on ROxford5K (Radenović et al., 2018a) for different values of output dimensions  $d$ , using features obtained from the pre-trained ResNet-101 of Revaud et al. (2019).

## D.1 TASKS AND DATASET

A summary of dataset statistics is available in Table C and examples for each dataset are available in Table D.

**MSMarco passages (Nguyen et al., 2016b):** question and answer dataset based on Bing queries. Very sparse annotation with a high number of false negatives. Queries (Q) and Documents (D) are from different domains due to size and content. Retrieval is asymmetric, because if you input a D as Q, the answer will not be D. Used only for pretraining as it has a set of training pairs for contrastive learning, while our aim is to perform self-supervision only. For this goal, we have chosen the other four datasets, that do not have a readily available set of training pairs (or triplets) for training, and thus self-supervision or unsupervised learning is required.

**ArguAna (Wachsmuth et al., 2018):** Counter-argument retrieval dataset. Queries and documents belong to the same domain, with some queries being a part of the corpus, which makes it not suitable for training on this dataset. Queries and documents come from the same domain in both size and content, however associated query-document pairs have inverse context (Q defends a point, D is a rebuttal of Q), so input Q should not retrieve Q, if it is on the database. Retrieval is asymmetric as a query should not retrieve itself.



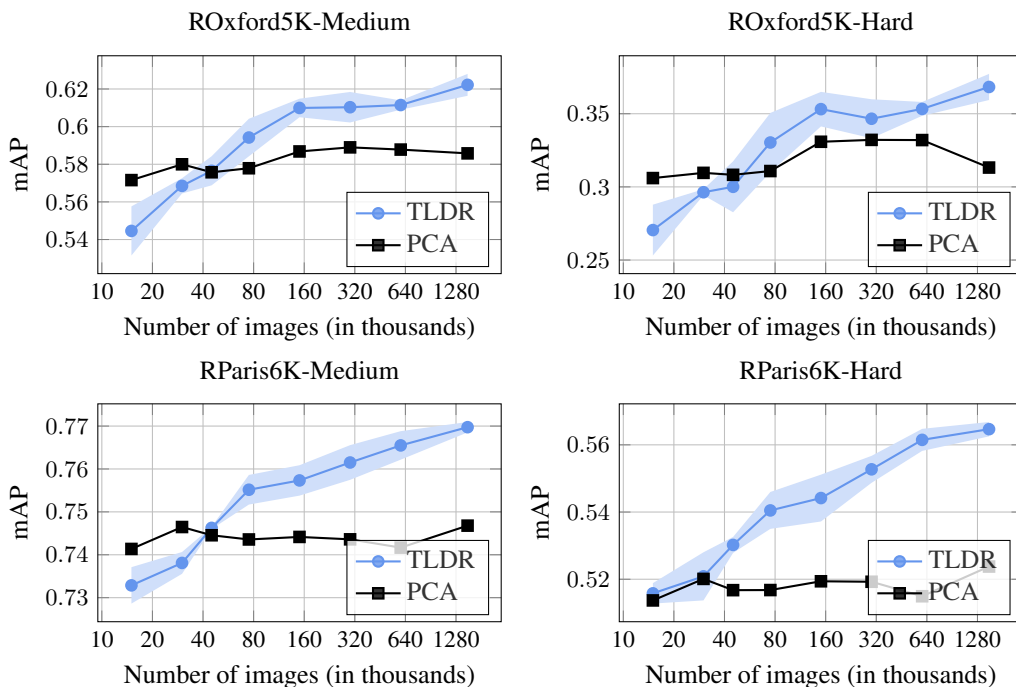


Figure D: **TLDR benefits from larger training sets.** Impact of the size of the training set on performance. TLDR uses a linear encoder and  $d = 128$ .




Figure E: **The surprising stability of TLDR across batch sizes.** Impact of the size of the training mini-batch on performance. TLDR uses a linear encoder and  $d = 128$ .

**Webis-Touché 2020 (Bondarenko et al., 2020; Wachsmuth et al., 2017):** Argument retrieval dataset. Queries and documents are from different domains due to size and content, with queries being questions and documents being support arguments for the question. Retrieval is asymmetric as a query should not retrieve itself.

**CQADupStack (Hoogveen et al., 2016):** Duplicate question retrieval from StackExchange subforums, composed of 12 different subforums. Corporuses are concatenated during training and mean result over all corporuses is used for testing (*i.e.* every corpus has equal weight even if the number of queries is different). Queries are titles of recent submissions, while documents are concatenation of titles and descriptions of existing ones. Queries and documents are from different domains due to size and content, with the query domain being a part of the document one (Queries are contained in the documents). Retrieval is symmetric as a query should return itself.

**Quora:** Duplicate question retrieval from the Quora platform. Queries are titles of recent submissions, while documents are titles of existing ones. Queries and documents are from the same domain concerning size and content. Retrieval is symmetric as a query should return itself.

Dataset	# Documents	# Queries	Avg positives per query	Avg query length	Avg document length	Retrieval type
Question answering (pretraining only)						
MSMarco	8.8M	6980	1.1	6	56	Asymmetric
Argument retrieval						
ArguANA	8674	1406	1	193	167	Asymmetric
Webis-Touché 2020	380k	49	49.2	7	292	Asymmetric
Duplicate question retrieval						
Quora	523k	5000	1.6	10	11	Symmetric
CQADupStack	457k	13145	1.4	9	129	Symmetric

Table C: Summary of **tasks and datasets** for the first stage document retrieval experiments presented in the Appendix.

Dataset	Query	Relevant-Documents
MSMARCO	what fruit is native to australia	<Paragraph> Passiflora herbertiana. A rare passion fruit native to Australia. Fruits are green-skinned, white fleshed, with an unknown edible rating. Some sources list the fruit as edible, sweet and tasty, while others list the fruits as being bitter and inedible. assiflora herbertiana. A rare passion fruit native to Australia...
ArguAna	Sexist advertising is subjective so would be too difficult to codify. Effective advertising appeals to the social, cultural, and personal values of consumers. Through the connection of values to products, services and ideas, advertising is able to accomplish its goal of adoption...	<Title> media modern culture television gender house would ban sexist advertising <Paragraph> Although there is a claim that sexist advertising is to difficult to codify, such codes have and are being developed to guide the advertising industry. These standards speak to advertising which demeans the status of women, objectifies them, and plays upon stereotypes about women which harm women and society in general. Earlier the Council of Europe was mentioned, Denmark, Norway and Australia as specific examples of codes or standards for evaluating sexist advertising which have been developed.
Touché-2020	Should the government allow illegal immigrants to become citizens?	<Title> America should support blanket amnesty for illegal immigrants. <Paragraph> Undocumented workers do not receive full Social Security benefits because they are not United States citizens " nor should they be until they seek citizenship legally. Illegal immigrants are legally obligated to pay taxes...
CQADupStack	Command to display first few and last few lines of a file	<Title> Combing head and tail in a single call via pipe <Paragraph> On a regular basis, I am piping the output of some program to either 'head' or 'tail'. Now, suppose that I want to see the first AND last 10 lines of piped output, such that I could do something like ./lotsoutput   headtail...
Quora	How long does it take to methamphetamine out of your blood?	<Paragraph> How long does it take the body to get rid of methamphetamine?

Table D: **Examples of queries and documents** from all the document retrieval datasets we use. Table extracted from [Thakur et al. \(2021\)](#); Note the difference of length between query and document in some datasets.

## D.2 EXPERIMENTAL RESULTS

We now test different combinations of the previously introduced datasets as dimensionality reduction and test datasets. The setup for all experiments is the same as the experiments in the main paper. In order to summarize our results (9 combinations of train/test datasets), we consider  $d = 64$  (second highest we investigate) as the comparison mark between PCA and TLDR. If TLDR outperforms PCA for all  $d \leq 64$ , we consider that it performed better than PCA, and otherwise we consider that PCA performed better than TLDR. Note that in all cases the lower the dimension the better TLDR performed against PCA. We also report which version of TLDR performed better,  $L > 0$  means that factorized linear is better than linear and  $L = 0$  the opposite. We present a summary of the experimental results in Table E, and provide depictions of some experiments in Figures F through I. From the results presented on the table, we derive two conclusions about TLDR:

- Differences in retrieval from pretraining to dimensionality reduction impacts results :** Looking into evaluation on argument retrieval, TLDR outperforms PCA. On the other hand, looking into evaluations on duplicate query retrieval, PCA is always able to outperform TLDR for  $d \geq 64$ . We infer that this must be derived from the difference in retrieval condition, as in all tests with asymmetric retrieval TLDR is able to outperform PCA. Note that symmetric retrieval and same domain for document and queries differs from the original pretraining task, and we posit that PCA is more robust to this type of change (which does not happen in our image retrieval experiments). Although we have this initial suspicion validated with 4 datasets a proper conclusion would need more dataset-pairs for experimentation, which we leave for future work.
- Choosing linear or factorized linear depends on the statistics of the dataset:** Analyzing the results we are able to detect that the choice of which version of TLDR one should use depends on the length of queries and documents of the original dataset. If both lengths are equal, factorized linear is better (ArguANA and Quora), if not then linear is the better

Dimensionality reduction	Argument Retrieval Duplicate Question	Test dataset			
		Argument retrieval		Duplicate query	
		Webis-Touché 2020	Webis-Touché 2020	Quora	CQADupStack
		Quora	TLDR (L>0)	TLDR (L=0)	PCA (L>0)
CQADupStack	TLDR (L>0)	TLDR (L=0)	PCA (L>0)	PCA (L=0)	

Table E: Summary of the results on document retrieval. (L=0) and (L>0) indicate which version of TLDR had better performance (linear and factorized linear respectively). Note that arguana is not suitable for training (not represented) and that we are not interested in using the same dataset for dimensionality reduction and test (thus the empty cells).

choice (Webis-Touché 2020 and CQADupStack). Even if by using ANCE representations we should not need to deal with these differences (we only tackle embeddings of fixed size), the statistics of the resulting embedding is different enough that it is detected by the batch normalization layer that is added for factorized linear.




Figure F: **Argument retrieval results on the ArguAna dataset using Webis-Touché 2020 for dimensionality reduction** for different values of output dimensions  $d$ . On the left we present Recall@100 and on the right we present NDCG@10.

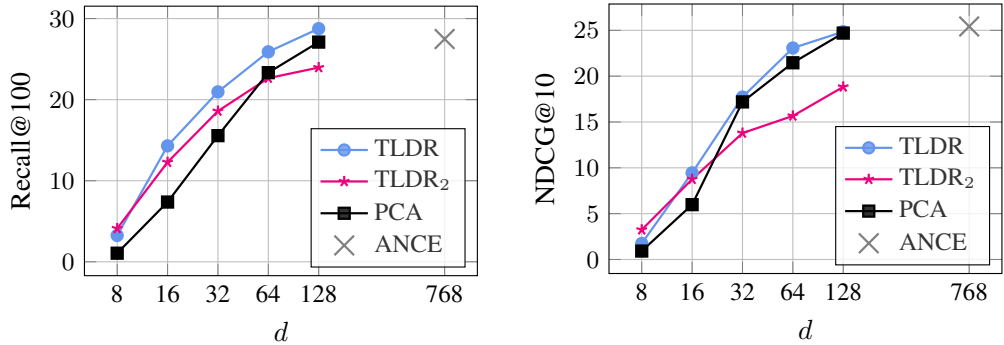


Figure G: **Argument retrieval results on the Webis-Touché 2020 dataset using Quora for dimensionality reduction** for different values of output dimensions  $d$ . On the left we present Recall@100 and on the right we present NDCG@10.

## E FASHIONMNIST: LEARNING FROM RAW PIXEL DATA AND VISUALIZATION

Although beyond the scope of what TLDR is designed for (see also discussion at the end of Section 1), in this section we present some basic results when using it for learning from raw pixel data and for visualizations, *i.e.* when reducing the output dimension to only  $d' = 2$ .

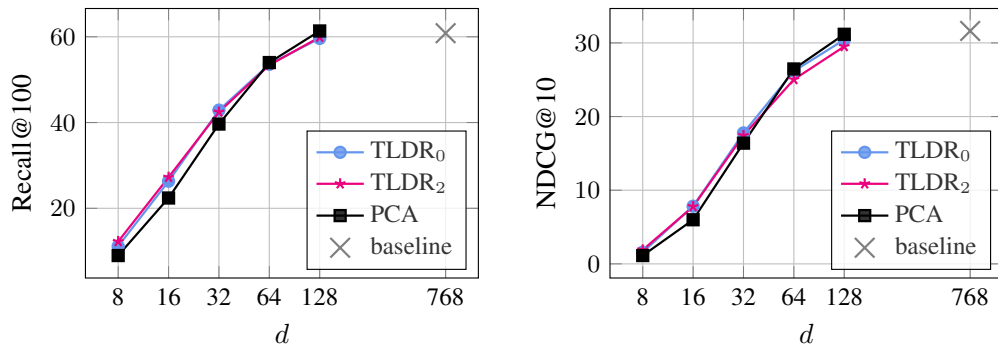


Figure H: **Duplicate question retrieval results on the CQADupstack dataset using Quora for dimensionality reduction** for different values of output dimensions  $d$ . On the left we present Recall@100 and on the right we present NDCG@10.




Figure I: **Duplicate question retrieval results on the Quora dataset using Webis-Touché 2020 for dimensionality reduction** for different values of output dimensions  $d$ . On the left we present Recall@100 and on the right we present NDCG@10.

**Learning from raw pixel data.** In Figure J we present results when learning directly from raw pixel data. We use the predefined splits and, following related work (McInnes et al., 2018), we measure and report accuracy after  $k$ -NN classifiers. We see that TLDR retains its gains over any other manifold learning method we tested. We have to note however that these results have to be taken with a pinch of salt, as a) the input pixel space is relatively simple compared to higher resolution natural images and b) to achieve such results we use the prior knowledge that we only have 10 classes and set high values for hyper-parameter  $k$ , *i.e.*  $k = 100$  for all methods compared.

**2D visualizations.** Let us first clarify that TLDR was not created with 2D outputs in mind; in fact, there are other excellent choices for visualization like  $t$ -SNE, UMAP (McInnes et al., 2018), TriMAP (Amid & Warmuth, 2019) or the recent Minimum-Distortion Embedding (MDE) (Agrawal et al., 2021) that we would use instead. In Figure K we show 2d visualizations when reducing the 60k training set of FashionMNIST to  $d = 2$  dimensions. We present results for TLDR,  $t$ -SNE (Van der Maaten & Hinton, 2008), UMAP (McInnes et al., 2018) and PyMDE (Agrawal et al., 2021). It is interesting how TLDR seems to be optimized for linear separability even for 2-dimensional outputs. For visualizations, we used the pyMDE library<sup>5</sup> provided by the authors of (Agrawal et al., 2021).

## F FURTHER DISCUSSIONS AND RELATED WORKS

**Graph diffusion for harder  $k$ -NN pairs.** Iscen et al. (2018b) improve the method presented in Hadsell et al. (2006) by mining harder positives and negative pairs for the contrastive loss via diffusion over the  $k$ -NN graph. Similar to Hadsell et al. (2006), they are interested in learning

<sup>5</sup><https://pymde.org/>




Figure J: **Results on the FashionMNIST dataset** as a function of the output dimensions  $d$ . We compare TLDR with PCA, PCA with whitening, UMAP and Isomap and report accuracy after  $k'$ -NN classifiers (with  $k' = 100$ ) following (McInnes et al., 2018). For TLDR and UMAP we set the number of neighbors  $k = 100$ . The performance of UMAP was very low for  $d' > 32$ .

(fine-tuning) the whole network and not just the dimensionality-reduction layer. Although it would be interesting to incorporate such ideas in TLDR, we consider it complementary and beyond the scope of this paper. Methods used for learning descriptor matching are also related; *e.g.* (Simonyan et al., 2014) formulates dimensionality reduction as a convex optimisation problem. Although the redundancy reduction objective can be formulated in many ways, *e.g.* via stochastic proximal gradient methods like Regularised Dual Averaging in (Simonyan et al., 2014), we believe that the *simplicity, immediacy and clarity* in which the Barlow Twins objective optimizes the output space is a strong advantage of TLDR.

**Graph diffusion for query expansion.** For the task of retrieval, assuming access to the search (test) database, methods like (Iscen et al., 2017; 2018a; Liu et al., 2019) utilize manifold learning on the the  $k$ -NN graph of the database to facilitate *query expansion*. We note that while these methods have shown great empirical performance on the same image retrieval datasets as we experiment on, we do not directly compare to them as their methodology and goals greatly differs from ours. We aim at being *invariant to the target dataset* (thus not performing learning on them), differently from the aforementioned methods they need access to the target dataset for learning, and to its  $k$ -NN graph during testing. TLDR is complementary to such graph diffusion techniques for query expansion.

**Relation to knowledge distillation.** Knowledge distillation (KD) (Hinton et al., 2015) aims at transferring knowledge from a pre-trained teacher network to a student one, often for neural network compression. One way to perform KD is relational KD (RKD) (Park et al., 2019; Tian et al., 2019; Lin et al., 2020), which transfers knowledge using relations between samples such as distance and angles. TLDR can be seen as a method for RKD. It enforces the student network (encoder) to reproduce a relational property (neighborhood) found on the teacher (the input space). However, there are some main differences to traditional distillation methods: i) the application: self-supervised retrieval instead of supervised classification (Hinton et al., 2015), contrastive (Tian et al., 2019; Lin et al., 2020) or self-supervision for classification (Fang et al., 2021), ii) the definition of the relations: abstract (neighbors), instead of measurable ones (distance, angle), which avoids normalization problems due to the dimensionality difference between teacher and student, and iii) the link between teacher and student: in our case, the student becomes a part of the teacher network at the end, instead of being a separate network.

**Relation to node embedding.** Node embedding methods aim at generating representations to graph nodes that are representative of the sample and its relations on the graph. In that sense, TLDR could be seen as learning embeddings for nodes on a graph. Compared to the traditional methods in this space, such as LINE (Tang et al., 2015), Node2Vec (Grover & Leskovec, 2016), DeepWalk (Perozzi et al., 2014), TLDR has three clear differences: i) does not rely on the edge strength; ii) regularization

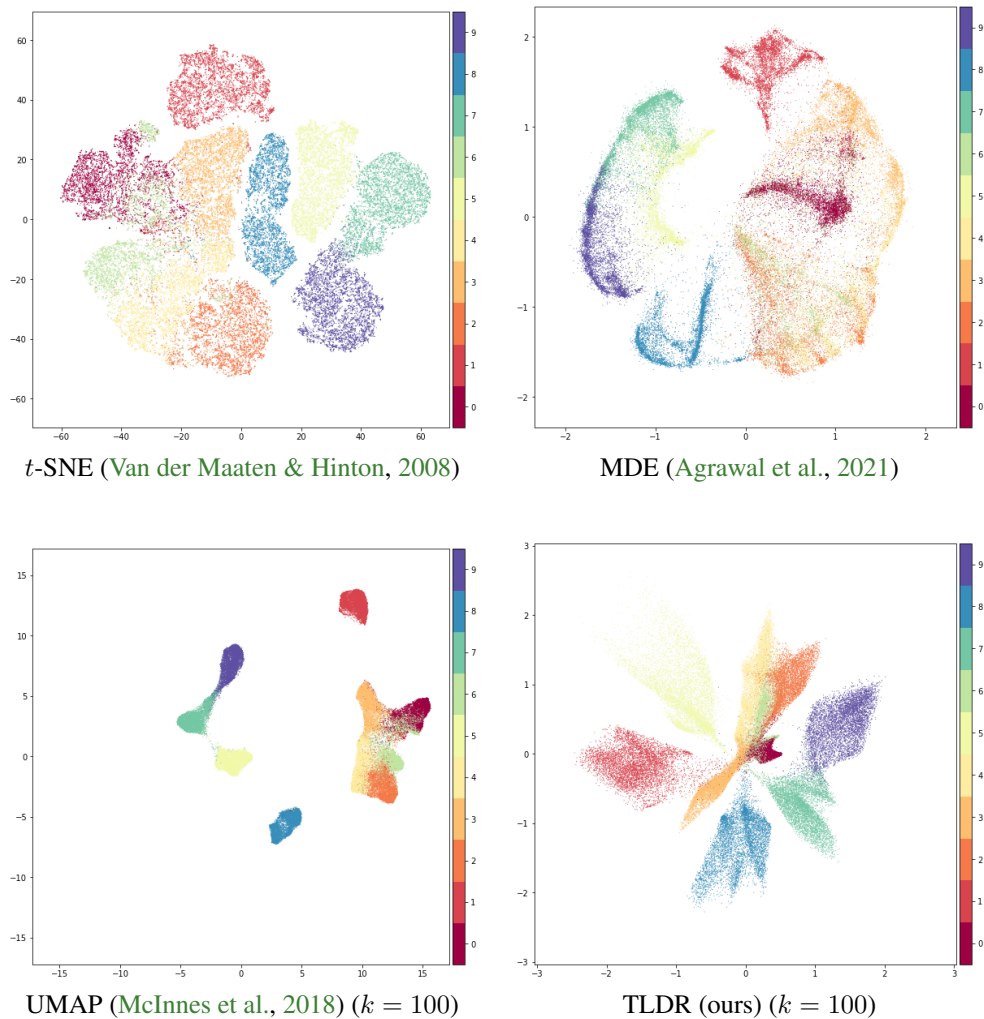


Figure K: **2D visualizations of the training set of FashionMNIST.** From top to bottom and left to right:  $t$ -SNE, MDE, UMAP and TLDR.

of the space based on the decorrelation of dimensions instead of L2-norm or orthogonality; and iii) only the 1-hop neighborhood information is used. More recent node embedding solutions are based on deep learning architectures that incorporate diffusion properties in the architecture like GCNs (Kipf & Welling, 2016; You et al., 2020), while TLDR achieves a similar effect via the Barlow Twins loss.

### F.1 LIMITATIONS OF OUR WORK

In the context of document retrieval we also tested TLDR on another task: duplicate question retrieval. In duplicate question retrieval TLDR was only able to outperform PCA for the lower dimension values ( $d=8,16,32$ ). We posit that TLDR does not achieve significant gains for the rest of the dimensions because duplicate task differs too much from the original pretraining task (QA on MSMarco dataset) in that the duplicate retrieval is symmetric (the documents retrieved by a query should also appear when we use the document as query), while pretraining and argument retrieval is assymmetric. In order to verify this, we performed ablations with different pairs of (dimensionality reduction, target dataset) and confirm that if the target dataset is a duplicate retrieval task TLDR is not able to outperform the compared method, but if we use duplicate retrieval only for dimensionality reduction and test on argument retrieval TLDR is able to outperform the compared methods. For full discussion and results *cf.* Section D.