

Image Hashing by Minimizing Independent Relaxed Wasserstein Distance

Khoa D. Doan^{1*}, Amir Kimiyaie¹, Saurav Manchanda² and Chandan K. Reddy¹

¹Virginia Tech

²University of Minnesota

khoadoan@vt.edu, am1rrad@vt.edu, manch043@umn.edu, reddy@cs.vt.edu

Abstract

Image hashing is a fundamental problem in the computer vision domain with various challenges, primarily, in terms of efficiency and effectiveness. Existing hashing methods lack a principled characterization of the goodness of the hash codes and a principled approach to learn the discrete hash functions that are being optimized in the continuous space. Adversarial autoencoders are shown to be able to implicitly learn a robust hash function that generates hash codes which are balanced and have low-quantization error. However, the existing adversarial autoencoders for hashing are too inefficient to be employed for large-scale image retrieval applications because of the minmax optimization procedure. In this paper, we propose an Independent Relaxed Wasserstein Autoencoder, which presents a novel, efficient hashing method that can implicitly learn the optimal hash function by directly training the adversarial autoencoder without any discriminator/critic. Our method is an order-of-magnitude more efficient and has a much lower sample complexity than the Optimal Transport formulation of the Wasserstein distance. The proposed method outperforms the current state-of-the-art image hashing methods for the retrieval task on several prominent image collections.

1 Introduction

The rapid growth of visual data, especially images, brings many challenges to the problem of finding similar items. Exact similarity search, which aims to exhaustively find all the relevant images, is often impractical due to its computational complexity. This is due to the fact that a complete linear scan of all the images in such massive databases is not feasible, especially when the database contains millions (or billions) of samples. Hashing is an approximate similarity search method which provides a principled approach for web-scale databases. In hashing, high-dimensional data

points are projected into a much smaller locality-preserving *binary* space via a hash function $f : x \rightarrow \{0, 1\}^m$. Approximate search for similar images can be efficiently conducted in this binary space using the Hamming distance calculation [Leskovec *et al.*, 2014]. Furthermore, the compact binary codes are storage-efficient.

The existing hashing methods can be broadly grouped into supervised and unsupervised hashing. Although supervised hashing offers a superior performance, unsupervised hashing is more suitable for large databases because it learns the hash function without any labeled data. A popular unsupervised hashing technique is Locality Sensitive Hashing (LSH) [Leskovec *et al.*, 2014]. The hashing problem has also been discussed in other papers, including the representative shallow hashing techniques such as Spectral Hashing (SpecHash) [Weiss *et al.*, 2009] and Iterative Quantization (ITQ) [Gong *et al.*, 2013], and deep hashing techniques such as SSDH [Yang *et al.*, 2018a].

Even though the existing methods have shown significant performance improvement in several image-hashing applications, they have two main drawbacks: (1) their objective functions are heuristically constructed without a principled characterization of the goodness of the hash codes, and (2) the gap of learning the discrete hash function in the continuous optimization framework is minimized heuristically with explicit constraints. The latter increases the time to tune the additional hyperparameters of the models. The recent work in [Doan *et al.*, 2019] shows that employing adversarial autoencoders for hashing avoids these explicitly constructed constraints. Their hashing model implicitly learns the optimal hash functions without the additional hyperparameter tuning. Furthermore, the adversarial autoencoders recover a better low-dimensional manifold of the data than the vanilla and variational autoencoders [Doan *et al.*, 2019; Makhzani *et al.*, 2015]. However, because they minimize the Jensen-Shannon (JS) divergence through the min-max optimization procedure, their model is not suitable for large-scale, real-world applications. On the other hand, replacing the JS divergence with Wasserstein distance has shown many advantages [Arjovsky *et al.*, 2017]. Moreover, it is possible to directly approximate the Wasserstein distance from the primal domain using the Optimal Transport (OT) formulation

*Contact Author

without the discriminator/critic [Iohara *et al.*, 2019]. However, solving the OT is computationally expensive and the Wasserstein estimate through the OT is known to have a high-variance [Iohara *et al.*, 2019].

To address these aforementioned challenges, we propose a novel unsupervised Independent Relaxed Wasserstein Autoencoder (IDR-AE) model for the image hashing problem. The proposed model implicitly learns the optimal hash function in a novel and efficient divergence minimization framework. The main contributions of the paper are as follows:

- Propose a principled hashing method that learns the locality-preserving hash codes without any explicitly defined additional constraints for the image hashing problem.
- Develop a novel, efficient optimization procedure that removes the discriminator/critic and the min-max optimization of the existing adversarial autoencoders. The proposed algorithm approximates the Wasserstein distance by a set of m independent one-dimensional projections. This algorithm is significantly more efficient than estimating the Wasserstein distance through the Optimal-Transport formulation.
- Demonstrate the superiority of the proposed model over the state-of-the-art hashing techniques on three prominent datasets; with both quantitative and qualitative performance analysis.

The rest of the paper is organized as follows. We discuss the related work in Section 2. In Section 3, we describe the details of our proposed method. Finally, we present quantitative and qualitative experimental results in Section 4 and conclude our discussion in Section 5.

2 Related Work

2.1 Image Hashing

Various unsupervised and supervised methods have been developed for hashing. Examples of supervised hashing methods include [Shen *et al.*, 2015; Yang *et al.*, 2018b; Ge *et al.*, 2014; Xia *et al.*, 2014; Cao *et al.*, 2018] and examples of unsupervised hashing methods include [Dizaji *et al.*, 2018; Huang *et al.*, 2017; Lin *et al.*, 2016; Huang *et al.*, 2016; Do *et al.*, 2016; He *et al.*, 2013; Heo *et al.*, 2012; Gong *et al.*, 2013; Weiss *et al.*, 2009; Salakhutdinov and Hinton, 2009]. While supervised methods demonstrate a superior performance over unsupervised ones, they require human-annotated datasets. Annotating massive-scale datasets, which are common in the image hashing domain, is an expensive and tedious task. Furthermore, besides the train/test distribution-mismatch problem, supervised methods easily get stuck in bad local optima when labeled data are limited. Thus, exploring the unsupervised hashing techniques is of great interest, especially in the image-hashing domain.

Hashing methods can also be categorized as either data independent or dependent. One of the most popular data-independent hashing technique is LSH [Leskovec *et al.*, 2014]. Data-dependent hashing includes popular methods such as SpecHash [Weiss *et al.*, 2009] and ITQ [Gong *et al.*, 2013]. Data-dependent hashing demonstrates a significant increase in retrieval performance because it considers the

data distribution. Hashing methods can also be categorized as shallow [Gionis *et al.*, 1999; Weiss *et al.*, 2009; Heo *et al.*, 2012] and deep hashing [Do *et al.*, 2016; Lin *et al.*, 2016; Huang *et al.*, 2017; Dizaji *et al.*, 2018]. The deep hashing methods can learn non-linear hash functions and have shown superiority over the shallow approaches.

In general, the existing hashing methods learn the optimal hash functions by minimizing the following training objective:

$$\min_f E_{x \sim D_x} L(x, f(x)) + E_{x \sim D_x} \sum_k \lambda_k \times H_k(f(x)) \quad (1)$$

where D_x is the data distribution, $L(x, f(x))$ is the locality-preserving loss of the hash function $f(x)$ and $H_k(f(x))$ is a hashing constraint with λ_k as its corresponding hyperparameter.

The existing hashing methods suffer from two limitations. First, in the existing methods, the loss function $L(x, f(x))$ is heuristically constructed. There is not a principled characterization of the locality-preserving property of the hash codes through the loss L . On the other hand, [Dai *et al.*, 2017] proves that learning to reconstruct the input in an autoencoder-based hashing model is exactly equivalent to preserving the original, input-space locality in the low-dimensional, hashing space. Second, in the relaxation from the discrete optimization to the continuous optimization, the constraints H_k 's are heuristically constructed. Examples of these constraints include bit-balance, bit-uncorrelation and low-quantization error [Doan *et al.*, 2019]. Because of the heuristic choices, the performance of the existing methods strongly depends on the choice of H_k 's and the adequate tuning of their "mixing" hyperparameters λ_k . In this paper, we leverage an adversarial autoencoder [Doan *et al.*, 2019], to implicitly learn the optimal hash functions without the inclusion of any H_k 's in the objective function.

2.2 Adversarial Learning

Generative Adversarial Network (GAN) has recently gained popularity due to its ability to generate realistic samples from the data distribution [Goodfellow *et al.*, 2014]. A prominent feature of GAN is its ability to "implicitly" match outputs of a deep network to a pre-defined distribution using the adversarial training procedure. Furthermore, adversarial learning has also been leveraged to regularize the latent space, as it helps in learning the intrinsic manifold of the data [Makhzani *et al.*, 2015; Doan *et al.*, 2019]. For example, the adversarially trained autoencoders can preserve manifold of the data in the low-dimensional latent space [Makhzani *et al.*, 2015]. Furthermore, regularizing the latent space with a "binary"-like distribution can guide the learning process to generate the optimal hash function without imposing any additional constraints H_k 's [Doan *et al.*, 2019]. However, training the adversarial autoencoders remains challenging and inefficient because of the alternating-optimization procedure (min-max game) between the generator and the discriminator. For example, [Doan *et al.*, 2019] employs the original non-saturating min-max GAN objective [Goodfellow *et al.*, 2014], which suffers from mode-collapse and vanishing gradient [Arjovsky *et al.*, 2017]. Moreover, in the min-max optimization, the generator's loss fluctuate during the

training instead of “descending”, making it extremely challenging to know when to stop the training process. Wasserstein GAN overcomes a few of these limitations (specifically, mode-collapse and vanishing gradient) [Arjovsky *et al.*, 2017]. However, it approximates the Wasserstein distance by employing the Kantorovich-Rubinstein dual, thus, resulting in a similar min-max game between the generator and the critic. [Tolstikhin *et al.*, 2017] proposes Wasserstein Autoencoder but still employs the min-max game. [Iohara *et al.*, 2019] directly estimates the Wasserstein distance by solving the the Optimal Transport (OT) problem. However, solving the OT has two main challenges. Firstly, its computation cost is $O(N^{2.5} \log(Nd))$ where N is the number of data points and d is the dimension of the data points. This is computationally expensive and can quickly become impractical as the number of training examples grows. Secondly, the OT-estimate of the Wasserstein distance has an exponential sample complexity [Deshpande *et al.*, 2019]. Intuitively, the learning process needs an exponential number of samples to reduce the variance in the estimated Wasserstein distances. On the other hand, Sliced Wasserstein Distance (SWD) approximates the Wasserstein distance by averaging many randomly projected one-dimensional Wasserstein distances [Deshpande *et al.*, 2018]. SWD has a polynomial sample complexity [Deshpande *et al.*, 2019]. However, in the high dimensional space, it becomes very likely that the random, one-dimensional projection does not lie on the manifold of the data. In other words, several random directions have near-zero one-dimensional Wasserstein distances. Therefore, in spite of the advantages over the OT-Wasserstein estimate, SWD needs a very large number of random directions in order to accurately estimate the Wasserstein distance.

In this paper, we address the limitations of these GAN-based approaches by robustly and efficiently minimizing a variant of the Wasserstein distance. Our proposed method does not employ any discriminator/critic and can estimate the Wasserstein distance with a polynomial sample complexity. Moreover, our method has a significant improvement in terms of its computational cost compared to the OT estimate.

3 Proposed Method

3.1 Problem statement

Given a data set $X = \{x^{(1)}, x^{(2)}, \dots, x^{(N)}\}$ of N images where $x \in R^d$, the goal of unsupervised hashing is to learn a hash function $f : x \rightarrow b$ that can generate binary hash code $b \in \{0, 1\}^m$ of the image. m denotes the length of the hash code b and it is typically much smaller than d .

3.2 Network architecture

We propose the IDR-AE network, which addresses the limitations of the previous application of the adversarial training for hashing in [Doan *et al.*, 2019]. Figure 1 shows the architecture of IDR-AE. The input image, represented by x_{pixel} , is input into the feature extractor $h : x_{pixel} \rightarrow x$, which is a convolutional neural network. We can either train h or employ a pre-trained feature extractor such as the VGG network.

The encoder, represented by the function $f : x \rightarrow b$, computes the low-dimensional representation b . Given the fea-




Figure 1: The network architecture of the proposed IDR-AE model.

ture vector x of an image, the output $b = f(x)$ is represented by the m independent probabilities $b_i = p(c_i = 1 | x, W_f)$, where W_f is the parameter of the encoder. To generate the hash codes, we simply compute $c_i = \mathbf{1}_{[b_i > 0.5]}$. The decoder, represented by the function $g : b \rightarrow x$, reconstructs the the input, denoted as \hat{x} .

In the following sections, we will discuss the motivations behind our proposed method, especially the novel adversarial loss.

3.3 Locality preservation of the hash codes

The autoencoder is trained to minimize the mean-squared error between the input and the reconstructed output, as below:

$$L_A = \|g(f(x)) - x\|_2^2 \quad (2)$$

[Dai *et al.*, 2017] shows that minimizing the reconstruction loss L_A is equivalent to preserving locality information of the data in the original input space. In other words, employing the autoencoder model for hashing produces a hash function that preserves that original input locality.

3.4 Implicit optimal hash function learning

Similar to the work in [Doan *et al.*, 2019], the encoder’s output b is regularized by an adversarial network. Specifically, we sample a vector z as the real data. In hashing, each component of z is independently and identically sampled from a one dimensional Bernoulli distribution with a parameter p . The sampling procedure defines a distribution P_z over z while the encoder defines a distribution P_b over the latent space b .

Instead of heuristically constructing the objective function with the constraints such as bit-balance, bit-uncorrelation and quantization error, [Doan *et al.*, 2019] implicitly learns the hash function satisfying these constraints by adversarially regularizing the autoencoder. Intuitively, in the existing hashing methods, the imposed constraints H_k ’s explicitly define the hash-code distribution over b . Instead, the adversarially trained autoencoder pre-defines the optimal hash-code distribution over z (the real data distribution) and implicitly aligns the encoder’s output distribution over b (the fake distribution) to the real distribution. One such optimal real distribution is the Bernoulli rdistribution with parameter $p = 0.5$ [Doan *et al.*, 2019].

In this work, we replace the Jensen-Shannon divergence with Wasserstein distance [Tolstikhin *et al.*, 2017]. However, we directly estimate the Wasserstein distance from its primal

domain by solving the OT problem. Specifically, in the primal domain, the Wasserstein distance is defined as follows:

$$W(P_b, P_z) = \inf_{\gamma \in \Pi(P_b, P_z)} \int_{(b,z) \sim \gamma} p(b,z) d(b,z) db dz \quad (3)$$

where $\Pi(P_b, P_z)$ is the set of all possible joint distributions of b and z whose marginals are P_b and P_z , respectively, and $d(b, z)$ is the cost of transporting one unit of mass from b to z . This approach is equivalent to solving the optimal transport (OT) problem, whose objective is to find the optimal transport plan to move masses from the generated distribution P_b to the true distribution P_z . Given two finite samples of N examples of b and N examples of z , the empirical transport cost can be formulated as the following Linear Programming (LP) problem:

$$\hat{W}(P_b, P_z) = \min_M \sum_i^N \sum_j^N M_{i,j} d(b^{(i)}, z^{(j)}) = \min_{W_f} M \odot D, \quad (4)$$

where M is the assignment matrix, D is the cost matrix where $D_{i,j} = d(b^{(i)}, z^{(j)})$ and \odot is the Hadamard product. The LP program has the following constraints:

$$\sum_i^N M_{i,j} = 1, \forall j = 1, \dots, N \quad (5)$$

$$\sum_j^N M_{i,j} = 1, \forall i = 1, \dots, N \quad (6)$$

$$M_{i,j} \in \{0, 1\}, \forall i = 1, \dots, N, \forall j = 1, \dots, N \quad (7)$$

It is important to note that both $M_{i,j}$ and $d(b^{(i)}, z^{(j)})$ are functions of W_f . We denote the “**distribution matching cost**” $\hat{W}(P_b, P_z)$ by L_G . Given the optimal assignment matrix M^* that solves this LP program, we can update the parameters of the generator, W_f , as follows:

$$W_f^{\text{new}} = W_f - \alpha \frac{\partial \hat{W}^*(P_b, P_z)}{\partial W_f} \quad (8)$$

$$\frac{\partial \hat{W}^*(P_b, P_z)}{\partial W_f} = \frac{\partial M^*}{\partial W_f} \odot D + M^* \odot \frac{\partial D}{\partial W_f} \quad (9)$$

While M^* depends on W_f , it is known that a small perturbation of W_f does not change the transport plan; that is $\frac{\partial M^*}{\partial W_f} = 0$. Thus, we update the generator using only the second term in Equation (9), $M^* \odot \frac{\partial D}{\partial W_f}$, which is easy to compute given a differentiable distance function $d(b, z)$. In our work, we use the Euclidean distance where $d(b, z) = \|b - z\|_2$.

We denote this algorithm **Optimal Transport Wasserstein AutoEncoder (LP-AE)**. The best method of solving the OT’s LP program has a cost of approximately $O(N^{2.5} \log(Nd))$ [Burkard *et al.*, 2009], where N is the

number of examples. While it is entirely possible to implement this LP program in a Stochastic Gradient Descent (SGD) training for small mini-batch sizes¹, it is computationally expensive for larger N .

3.5 Independent Relaxed Wasserstein Distance

In our experiments, we observe that it is difficult to match P_b to P_z by solving the OT. We conjecture that the reason is because OT is a poor estimate of the Wasserstein distance. The OT estimate of the Wasserstein distance has high variance when the number of samples N in the mini-batches is small [Iohara *et al.*, 2019]. On the other hand, each bit of b can be considered as a **one-dimensional projection of the data** that lies in the hashing space. Therefore, this motivates us to estimate the Wasserstein distance, as follows:

$$L_G = \frac{1}{m} \sum_i^m W(P_{b_i}, P_{z_i}) \quad (10)$$

where $W(P_{b_i}, P_{z_i})$ is the one-dimensional Wasserstein distance on the one-dimensional projection of bit i . Solving the OT in the one-dimensional space has a significantly small computational cost [Deshpande *et al.*, 2018]. The cost of such operation is equivalent to the one-dimensional array sort plus the distance calculation, or $O(n \log(n + d))$. This is an order of magnitude faster than the OT’s computational cost of $O(N^{2.5} \log(Nd))$. We call this algorithm **InD**ependent **R**elaxed **W**asserstein **A**uto**E**ncoder (IDR-AE).

Our proposed estimate is similar to SWD with one importance distinction. SWD employs a large number of random directions in order to accurately estimate the distance. Our proposed formulation estimates the Wasserstein distance from the directions that best separate the data generated b and the real data z . However, similar to SWD, estimating the Wasserstein distance from the m one dimensional projections has a polynomial sample complexity [Deshpande *et al.*, 2019]. This is an important advantage over the OT estimation, which has an exponential sample complexity. In our experiments, the OT approach frequently fails to match the generated hash code distribution to the real distribution z .

The objective function of the IDR-AE can be written as follows:

$$L = L_A + \lambda \times L_G \quad (11)$$

In our experiments, we observe that setting $\lambda = 0.01$ consistently achieves the best performance.

4 Experiments

4.1 Datasets Used

We utilize the following datasets in our performance evaluation experiments:

- **CIFAR10** [Krizhevsky *et al.*, 2009]: a dataset of 60,000 natural images categorized uniformly into 10 labels. We randomly select 1000 images from each label for the query set and use the remaining images for the training and retrieval sets. Therefore, the query set contains

¹<https://github.com/gatagat/lap>

Table 1: Performance comparison of different methods using P@1000. The best P@1000 value for each experiment is in bold.

Method	CIFAR10			FLICKR25K			PLACE365		
	32 bits	64 bits	128 bits	32 bits	64 bits	128 bits	32 bits	64 bits	128 bits
LSH	0.1721	0.1731	0.2165	0.0131	0.0188	0.0228	0.0058	0.0101	0.0081
SpecHash	0.1995	0.1982	0.1930	0.0213	0.0233	0.0229	0.0075	0.0076	0.0074
ITQ	0.2424	0.2585	0.2700	0.0264	0.0302	0.0305	0.0088	0.0096	0.0101
SGH	0.1672	0.1803	0.1888	0.0130	0.0137	0.0145	0.0061	0.0061	0.0137
SSDH	0.2218	0.1766	0.1766	0.0260	0.0271	0.0282	0.0149	0.0160	0.0282
WGAN-AE	0.1948	0.2170	0.2284	0.0219	0.0269	0.0289	0.0158	0.0174	0.0181
LP-AE	0.2370	0.2406	0.2396	0.0222	0.0273	0.0293	0.0147	0.0189	0.0240
IDR-AE	0.2692	0.2747	0.2961	0.0282	0.0336	0.0365	0.0205	0.0259	0.0301

10,000 images and the training/retrieval set contains the same 50,000 images.

- **FLICKR25K** [Bui *et al.*, 2018]: a dataset of 25,000 social photographic images downloaded from Flickr². The dataset is organized into 250 labels. We randomly select 20 images from each label for the query set and similarly use the remaining images for the training and retrieval sets. The final query dataset contains 5,000 images and the training/retrieval set contains the same 20,000 images.
- **PLACE365**³: a dataset of 1.8 millions of scenery images organized into 365 labels. We randomly select 10 images from each label for the query set and 500 images from each label for the training and retrieval sets. The final query dataset contains 3,650 images and the training/retrieval set contains 182,500 images.

4.2 Evaluation Metrics

For evaluating the performance of the proposed model, we follow the standard evaluation mechanism that is widely accepted in the context of image hashing - the **precision@R (P@R)** and **mean average precision (MAP)**. Given the query images, MAP is calculated as follows:

$$AP(q) = \frac{1}{N_q} \sum_{r=1}^N P@r \times \delta(r) \quad (12)$$

$$MAP = \frac{1}{Q} \sum_{q=1}^Q AP(q), \quad (13)$$

where N is the size of the retrieval set, N_q is the number of all relevant images in this set, Q is the size of the query set and $\delta(r) = 1$ only when the r -th retrieved image is relevant to the query image; otherwise $\delta(r) = 0$. A retrieved image is relevant if its ground-truth label is the same as the label of the query image.

4.3 Compared Methods

We compare the performance of the proposed method with various representative unsupervised image hashing methods.

²<https://www.flickr.com/>

³<http://places2.csail.mit.edu/>

- **Locality Sensitive Hashing (LSH)** [Leskovec *et al.*, 2014]: the popular data-independent, shallow hashing method using random projection.
- **Spectral Hashing (SpecHash)** [Weiss *et al.*, 2009]: an unsupervised shallow hashing method whose goals are to preserve locality and find balanced, uncorrelated hashes by solving the Eigenvector problem.
- **Iterative Quantization (ITQ)** [Gong *et al.*, 2013]: the state-of-the-art shallow hashing method that alternately minimizes the quantization error to achieve better hash codes.
- **Stochastic Generative Hashing (SGH)** [Dai *et al.*, 2017]: a representative hashing method that, similar to our proposed method, also minimizes the reconstruction loss in an autoencoder model.
- **Semantic Structure-based Deep Hashing (SSDH)** [Yang *et al.*, 2018a]: the state-of-the-art unsupervised deep hashing method that learns that hash function by preserving the heuristically-defined semantic structure of the data.
- **Wasserstein Adversarial Autoencoder (WGAN-AE)**: the adversarial autoencoder model for hashing which employs the critic that estimates the Wasserstein from the dual domain.
- **OT-Wasserstein Adversarial Autoencoder (LP-AE)**: the adversarial autoencoder model for hashing which directly minimizes the Wasserstein distance using the OT formulation in the primal domain.
- **Our proposed method (IDR-AE)**: our proposed adversarial autoencoder model for hashing.

Implementation Details

For the deep methods, we use raw images as input. For the existing shallow hashing techniques, we employ a pre-trained VGG network and extract the corresponding $fc7$ feature vectors of the images [Simonyan and Zisserman, 2014]. For our model, we employ VGG for the feature extractor h . The encoder/decoder are multi-layer perceptrons (MLP). We implement our proposed method in pyTorch⁴ and train our model using Stochastic Gradient Descent (SGD) along with Adam optimizer [Kingma and Ba, 2014]. We use a mini-batch size of 128 examples for IDR-AE and WGAN-AE. For LP-AE, we try different mini-batch sizes ranging from 128 to 512.

⁴<http://pytorch.org/>

Table 2: Performance comparison of different methods using MAP for $m = 64$ bits. The best MAP values are shown in bold.

Method	CIFAR10	FLICKR25K	PLACE365
LSH	0.1477	0.0348	0.0101
SpecHash	0.1265	0.0535	0.0121
ITQ	0.1824	0.0535	0.0163
SGH	0.1372	0.0196	0.0052
SSDH	0.1664	0.0504	0.0107
WGAN-AE	0.1775	0.0642	0.0161
LP-AE	0.1777	0.0658	0.0165
IDR-AE	0.2084	0.0790	0.0184

For each method, we perform a grid search to find the best hyper-parameters. Then, we evaluate the methods with their best configurations on three different samples of the data and average the metrics across three runs.

4.4 Performance Results

In this experiment, we measure the performance of compared methods in the image retrieval task. Table 1 shows the P@1000 results across different lengths of the hash codes. IDR-AE consistently outperforms all the baseline methods at different lengths of the hash codes. Specifically, IDR-AE has a relative performance improvement of at least 10% in CIFAR10 and FLICKR25K. Similarly, in Table 2, we report the MAP results for all the methods at different lengths of the hash codes. Again, IDR-AE consistently achieves the best MAP results. The improvements of our proposed method over the baselines are statistically significant according to the corresponding paired t-tests (p -value < 0.01).

The P@1000 and MAP results demonstrate the superiority of our proposed method in the image hashing problem.

4.5 Ablation Study

We further evaluate the effectiveness of the proposed adversarial learning procedure using an ablation study. Figure 2 shows the Precision-Recall curves of the different Autoencoder-based hashing models. AE denotes the vanilla Autoencoder without any adversarial regularization. We observe that all adversarial-based Autoencoder models outperform AE. This demonstrates the importance of the adversarial learning for Autoencoders in image hashing. Furthermore, replacing the dual Wasserstein estimate (in WGAN-AE) and the OT estimate (in LP-AE) with our estimate further improves the retrieval performance.

4.6 The Effectiveness of Adversarial Training

In this experiment, we compare the training time, as well as the behavior of L_G during training, of the various adversarial learning procedures which are discussed in this paper. The training time of WGAN-AE, LP-AE and IDR-AE are shown in Figure 3 for various datasets. We train CIFAR10 and FLICKR25K datasets for 100 epochs and PLACE364 for 50 epochs. In Figure 3, the training time of IDR-AE is significantly reduced compared to the training of LP-AE. In Figure 4, we observe that the training loss of L_G for WGAN-AE fluctuates during training while it decreases for IDR-AE. With the advantages of the single-optimization objective and

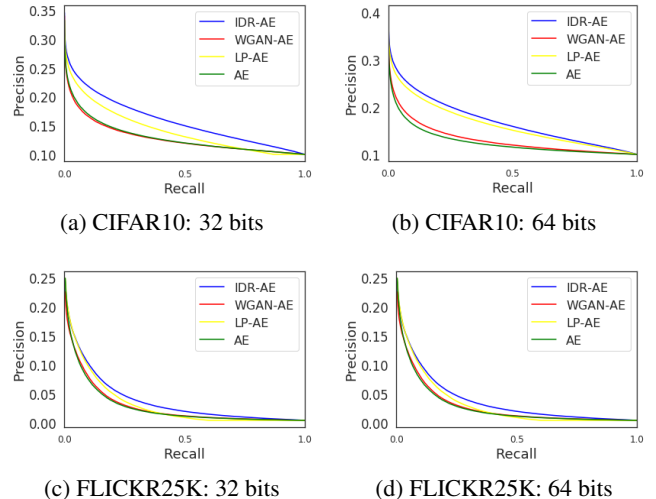


Figure 2: Precision-Recall curves of code lengths of 32 and 64 bits.

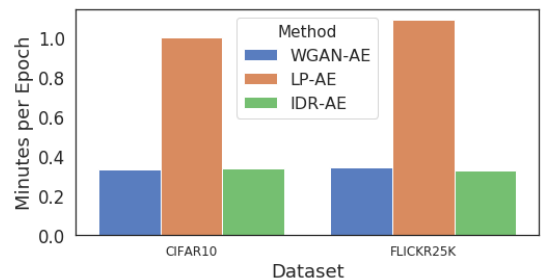


Figure 3: Training time per Epoch on CIFAR10 and FLICKR25k datasets (using a mini-batch size of 128 examples).

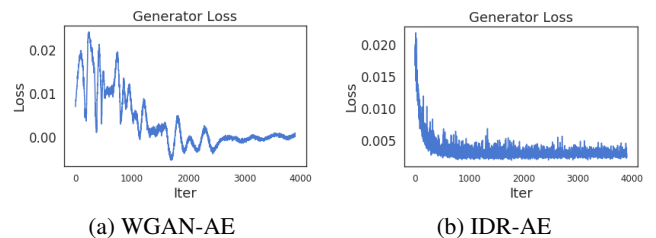


Figure 4: The loss L_G during training.

the efficient training time, IDR-AE can be employed in any real-world image-hashing applications.

5 Conclusion

We proposed a novel adversarial autoencoder model for the image hashing problem. To achieve this, we developed a new and efficient mechanism to train the adversarial autoencoders without the critics. Our model learns hash codes that preserves the locality information in the original data. Our model trains significantly faster than the OT-based adversarial autoencoder. We proved that our proposed hashing model

outperforms all existing state-of-the-art image hashing methods. Our work makes a step towards leveraging an efficient, robust adversarial autoencoder for the image hashing problem and we envision that our work will serve as a motivation for improving other adversarially trained models by exploring the primal domain of the Wasserstein distances. Our code will be available on GitHub for reproducibility.

References

- [Arjovsky *et al.*, 2017] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [Bui *et al.*, 2018] Tu Bui, Leonardo Ribeiro, Moacir Ponti, and John Collomosse. Sketching out the details: Sketch-based image retrieval using convolutional neural networks with multi-stage regression. *Computers & Graphics*, 71:77–87, 2018.
- [Burkard *et al.*, 2009] Rainer E Burkard, Mauro Dell’Amico, and Silvano Martello. *Assignment problems*. Springer, 2009.
- [Cao *et al.*, 2018] Yue Cao, Bin Liu, Mingsheng Long, and Jianmin Wang. Hashgan: Deep learning to hash with pair conditional wasserstein gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1287–1296, 2018.
- [Dai *et al.*, 2017] Bo Dai, Ruiqi Guo, Sanjiv Kumar, Niao He, and Le Song. Stochastic generative hashing. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 913–922. JMLR. org, 2017.
- [Deshpande *et al.*, 2018] Ishan Deshpande, Ziyu Zhang, and Alexander G Schwing. Generative modeling using the sliced wasserstein distance. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3483–3491, 2018.
- [Deshpande *et al.*, 2019] Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David Forsyth, and Alexander G Schwing. Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10648–10656, 2019.
- [Dizaji *et al.*, 2018] Kamran Ghasedi Dizaji, Feng Zheng, Najmeh Sadoughi Nourabadi, Yanhua Yang, Cheng Deng, and Heng Huang. Unsupervised deep generative adversarial hashing network. In *CVPR 2018*, 2018.
- [Do *et al.*, 2016] Thanh-Toan Do, Anh-Dzung Doan, and Ngai-Man Cheung. Learning to hash with binary deep neural network. In *European Conference on Computer Vision*, pages 219–234. Springer, 2016.
- [Doan *et al.*, 2019] Khoa D Doan, Pranjul Yadav, and Chandan K Reddy. Adversarial factorization autoencoder for look-alike modeling. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2803–2812, 2019.
- [Ge *et al.*, 2014] Tiezheng Ge, Kaiming He, and Jian Sun. Graph cuts for supervised binary coding. In *European Conference on Computer Vision*, pages 250–264. Springer, 2014.
- [Gionis *et al.*, 1999] Aristides Gionis, Piotr Indyk, Rajeev Motwani, et al. Similarity search in high dimensions via hashing. In *Vldb*, volume 99, pages 518–529, 1999.
- [Gong *et al.*, 2013] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2916–2929, 2013.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [He *et al.*, 2013] Kaiming He, Fang Wen, and Jian Sun. K-means hashing: An affinity-preserving quantization method for learning binary compact codes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2938–2945, 2013.
- [Heo *et al.*, 2012] Jae-Pil Heo, Youngwoon Lee, Junfeng He, Shih-Fu Chang, and Sung-Eui Yoon. Spherical hashing. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2957–2964. IEEE, 2012.
- [Huang *et al.*, 2016] Chen Huang, Chen Change Loy, and Xiaoou Tang. Unsupervised learning of discriminative attributes and visual representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5175–5184, 2016.
- [Huang *et al.*, 2017] Shanshan Huang, Yichao Xiong, Ya Zhang, and Jia Wang. Unsupervised triplet hashing for fast image retrieval. In *Proceedings of the on Thematic Workshops of ACM Multimedia 2017*, pages 84–92. ACM, 2017.
- [Iohara *et al.*, 2019] Akihiro Iohara, Takahito Ogawa, and Toshiyuki Tanaka. Generative model based on minimizing exact empirical wasserstein distance, 2019.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Leskovec *et al.*, 2014] Jure Leskovec, Anand Rajaraman, and Jeffrey David Ullman. *Mining of massive datasets*. Cambridge university press, 2014.
- [Lin *et al.*, 2016] Kevin Lin, Jiwen Lu, Chu-Song Chen, and Jie Zhou. Learning compact binary descriptors with unsupervised deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1183–1192, 2016.
- [Makhzani *et al.*, 2015] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.

- [Salakhutdinov and Hinton, 2009] Ruslan Salakhutdinov and Geoffrey Hinton. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7):969–978, 2009.
- [Shen *et al.*, 2015] Fumin Shen, Chunhua Shen, Wei Liu, and Heng Tao Shen. Supervised discrete hashing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 37–45, 2015.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Tolstikhin *et al.*, 2017] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
- [Weiss *et al.*, 2009] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *Advances in neural information processing systems*, pages 1753–1760, 2009.
- [Xia *et al.*, 2014] Rongkai Xia, Yan Pan, Hanjiang Lai, Cong Liu, and Shuicheng Yan. Supervised hashing for image retrieval via image representation learning. In *AAAI*, volume 1, page 2, 2014.
- [Yang *et al.*, 2018a] Erkun Yang, Cheng Deng, Tongliang Liu, Wei Liu, and Dacheng Tao. Semantic structure-based unsupervised deep hashing. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 1064–1070, 2018.
- [Yang *et al.*, 2018b] Huei-Fang Yang, Kevin Lin, and Chu-Song Chen. Supervised learning of semantics-preserving hash via deep convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 40(2):437–451, 2018.