

Discrete Infomax Codes for Supervised Representation Learning

Yoonho Lee ^{1,*} , Wonjae Kim ² , Wonpyo Park ³  and Seungjin Choi ⁴ ¹ Stanford AI Lab, Stanford University, Stanford, CA 94305, USA² NAVER AI Lab, Seongnam 13561, Korea; wonjae.kim@navercorp.com³ Standigm, Seoul 06234, Korea; wppark.pio@gmail.com⁴ Intellicode & BARO AI Academy, Seoul 06367, Korea; seungjin.choi.mlg@gmail.com

* Correspondence: yoonho@cs.stanford.edu

Abstract: For high-dimensional data such as images, learning an encoder that can output a compact yet informative representation is a key task on its own, in addition to facilitating subsequent processing of data. We present a model that produces discrete infomax codes (DIMCO); we train a probabilistic encoder that yields k -way d -dimensional codes associated with input data. Our model maximizes the mutual information between codes and ground-truth class labels, with a regularization which encourages entries of a codeword to be statistically independent. In this context, we show that the infomax principle also justifies existing loss functions, such as cross-entropy as its special cases. Our analysis also shows that using shorter codes reduces overfitting in the context of few-shot classification, and our various experiments show this implicit task-level regularization effect of DIMCO. Furthermore, we show that the codes learned by DIMCO are efficient in terms of both memory and retrieval time compared to prior methods.

Keywords: infomax; discrete codes; representation learning; few-shot classification



Citation: Lee, Y.; Kim, W.; Park, W.; Choi, S. Discrete Infomax Codes for Supervised Representation Learning. *Entropy* **2022**, *24*, 501. <https://doi.org/10.3390/e24040501>

Academic Editors: Sergio Cruces, Iván Durán-Díaz, Rubén Martín-Clemente and Andrzej Cichocki

Received: 26 February 2022

Accepted: 29 March 2022

Published: 2 April 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Metric learning and few-shot classification are two problem settings that test a model's ability to classify data from classes that were unseen during training. Such problems are also commonly interpreted as testing meta-learning ability, since the process of constructing a classifier with examples from new classes can be seen as learning. Many recent works [1–4] tackled this problem by learning a continuous embedding ($\tilde{x} \in \mathbb{R}^n$) of datapoints. Such models compare pairs of embeddings using, e.g., Euclidean distance to perform nearest neighbor classification. However, it remains unclear whether such models effectively utilize the entire space of \mathbb{R}^n .

Information theory provides a framework for effectively asking such questions about representation schemes. In particular, the information bottleneck principle [5,6] characterizes the optimality of a representation. This principle states that the optimal representation \tilde{X} is one that maximally compresses the input X while also being predictive of labels Y . From this viewpoint, we see that the previous methods which map data to \mathbb{R}^n focus on being predictive of labels Y without considering the compression of X .

The degree of compression of an embedding is the number of bits it reflects about the original data. Note that for continuous embeddings, each of the n numbers in a n -dimensional embedding requires 32 bits. It is unlikely that unconstrained optimization of such embeddings use all of these $32n$ bits effectively. We propose to resolve this limitation by instead using *discrete embeddings* and controlling the number of bits in each dimension via hyperparameters. To this end, we propose a model that produces discrete infomax codes (DIMCO) via an end-to-end learnable neural network encoder.

This work's primary contributions are as follows. We use mutual information as an objective for learning embeddings, and propose an efficient method of estimating it in the discrete case. We experimentally demonstrate that learned discrete embeddings

are more memory and time-efficient compared to continuous embeddings. Our experiments also show that using discrete embeddings help meta-generalization by acting as an information bottleneck. We also provide theoretical support for this connection through an information-theoretic probably approximately correct (PAC) bound that shows the generalization properties of learned discrete codes.

This paper is organized as follows. We propose our model for learning discrete codes in Section 2. We justify our loss function and also provide a generalization bound for our setup in Section 3. We compare our method to related work in Section 4, and present experimental results in Section 5. Finally, we conclude our paper in Section 6.

2. Discrete Infomax Codes (Dimco)

We present our model which produces discrete infomax codes (DIMCO). A deep neural network is trained end-to-end to learn k -way d -dimensional discrete codes that maximally preserve the information on labels. We outline the training procedure in Algorithm 1, and we also illustrate the overall structure in the case of 4-way 3-dimensional codes ($k = 4, d = 3$) in Figure 1.

Algorithm 1 DIMCO training procedure.

```

Initialize network parameters  $\theta$ 
repeat
  Sample a batch  $\{(\mathbf{x}^n, y^n)\}$ 
  Compute logits  $\{l_{i,j}^n\} = f(\mathbf{x}^n; \theta)$  via (1)
  Compute probs  $\{p_{i,j}^n\} = \text{softmax}(\{l_{i,j}^n\})$  via (2)
  Update  $\theta$  to minimize loss  $\mathcal{L}(\{p_{i,j}^n\}, \{y^n\})$  via (10)
until converged
  
```

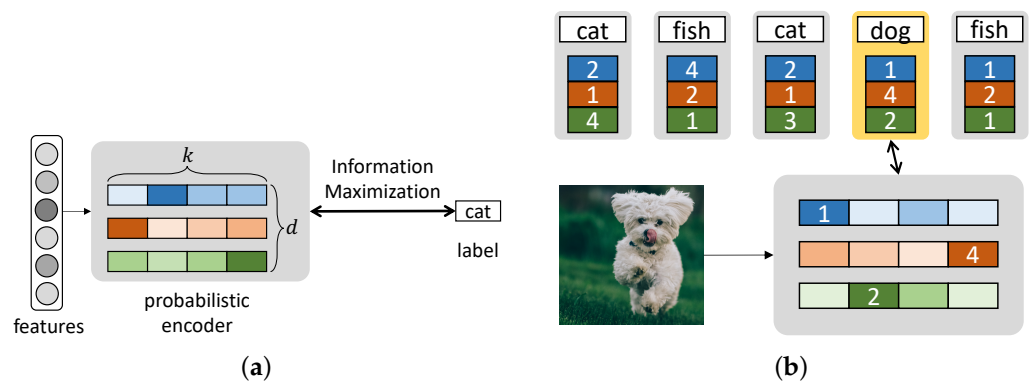


Figure 1. A graphical overview of discrete infomax codes (DIMCO). (a) Discrete codes are produced by a probabilistic encoder that maps each datapoint to a distribution over k -way d -dimensional discrete codes. The encoder is trained to maximize the mutual information between code distribution and label distribution. (b) Given a query image, we compare it against a support set of discrete codes and corresponding labels. Our similarity metric is the query’s log probability for each discrete code.

2.1. Learnable Discrete Codes

Suppose that we are given a set of labeled examples, which are realizations of random variables $(X, Y) \sim p(\mathbf{x}, y)$, where X is the continuous input, and its corresponding discrete label is Y . Realizations of X and Y are denoted by $\mathbf{x} \in \mathbb{R}^D$ and $y \in \{1, \dots, c\}$. The codebook \tilde{X} serves as a compressed representation of X .

We constructed a probabilistic encoder $p(\tilde{\mathbf{x}}|\mathbf{x})$ —which is implemented by a deep neural network—that maps an input \mathbf{x} to a k -way d -dimensional code $\tilde{\mathbf{x}} \in \{1, 2, \dots, k\}^d$. That is, each entry of $\tilde{\mathbf{x}}$ takes on one of k possible values, and the cardinality of \tilde{X} is $|\tilde{X}| = k^d$. Special cases of this coding scheme include k -way class labels ($d = 1$), d -dimensional binary codes ($k = 2$), and even fixed-length decimal integers ($k = 10$).

We now describe our model which produces discrete infomax codes. A neural network encoder $f(\mathbf{x}; \theta)$ outputs k -dimensional categorical distributions, $\text{Cat}(p_{i,1}, \dots, p_{i,k})$. Here, $p_{i,j}(\mathbf{x})$ represents the probability that output variable i takes on value j , consuming \mathbf{x} as an input, for $i = 1, \dots, d$ and $j = 1, \dots, k$. The encoder takes \mathbf{x} as an input to produce logits $l_{i,j} = f(x)_{i,j}$, which form a $d \times k$ matrix:

$$\begin{bmatrix} l_{1,1} & l_{1,2} & l_{1,3} & \dots & l_{1,k} \\ l_{2,1} & l_{2,2} & l_{2,3} & \dots & l_{2,k} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ l_{d,1} & l_{d,2} & l_{d,3} & \dots & l_{d,k} \end{bmatrix}. \quad (1)$$

These logits undergo softmax functions to yield

$$p_{i,j} = \frac{\exp(l_{i,j})}{\sum_{j=1}^k \exp(l_{i,j})}. \quad (2)$$

Each example \mathbf{x} in the training set is assigned a codeword $\tilde{\mathbf{x}} = [\tilde{x}_1, \dots, \tilde{x}_d]^\top$, each entry of which is determined by one of k events that is most probable; i.e.,

$$\tilde{\mathbf{x}} = \left[\overbrace{\arg \max_j p_{1j}}^{\tilde{x}_1}, \dots, \overbrace{\arg \max_j p_{dj}}^{\tilde{x}_d} \right]^\top. \quad (3)$$

While the stochastic encoder $p(\tilde{\mathbf{x}}|\mathbf{x})$ induces a soft partitioning of input data, codewords assigned by the rule in (3) yield a hard partitioning of X .

2.2. Loss Function

The i -th symbol is assumed to be sampled from the resulting categorical distribution $\text{Cat}(p_{i1}, \dots, p_{ik})$. We denote the resulting distribution over codes as \tilde{X} and a code as $\tilde{\mathbf{x}}$. Instead of sampling $\tilde{\mathbf{x}} \sim \tilde{X}$ during training, we use a loss function that optimizes the expected performance of the entire distribution \tilde{X} .

We train the encoder by maximizing the mutual information between the distributions of codes \tilde{X} and labels Y . The mutual information is a symmetric quantity that measures the amount of information shared between two random variables. It is defined as

$$I(\tilde{X}; Y) = H(\tilde{X}) - H(\tilde{X}|Y). \quad (4)$$

Since \tilde{X} and Y are discrete, their mutual information is bounded from both above and below as $0 \leq I(\tilde{X}; Y) \leq \log |\tilde{X}| = d \log k$. To optimize the mutual information, the encoder directly computes empirical estimates of the two terms on the right-hand side of (4). Note that both terms consist of entropies of categorical distributions, which have the general closed-form formula:

$$H(p_1, \dots, p_n) = \sum_{i=1}^n p_i \log p_i. \quad (5)$$

Let \bar{p}_{ij} be the empirical average of p_{ij} calculated using data points in a batch. Then, \bar{p}_{ij} is an empirical estimate of the marginal distribution $p(\tilde{\mathbf{x}})$. We compute the empirical estimate of $H(\tilde{X})$ by adding its entropy estimate for each dimension.

$$\hat{H}(\tilde{X}) = \sum_{i=1}^d H(\bar{p}_{i1}, \dots, \bar{p}_{ik}). \quad (6)$$

We can also compute

$$\hat{H}(\tilde{X}|Y) = \sum_{y=1}^c p(Y=y) \hat{H}(\tilde{X}|Y=y), \quad (7)$$

where c is the number of classes. The marginal probability $p(Y=y)$ is the frequency of class y in the minibatch, and $H(\tilde{X}|Y=y)$ can be computed by computing (6) using only datapoints which belong to class y . We emphasize that such a closed-form estimation of $I(\tilde{X}; Y)$ is only possible because we are using discrete codes. If \tilde{X} were instead a continuous variable, we would only be able to maximize an approximation of $I(\tilde{X}; Y)$ (e.g., Belghazi et al. [7]).

We briefly examine the loss function (4) to see why maximizing it results in discriminative \tilde{X} . Maximizing $H(\tilde{X})$ encourages the distribution of all codes to be as dispersed as possible, and minimizing $H(\tilde{X}|Y)$ encourages the average embedding of each class to be as concentrated as possible. Thus, the overall loss $I(\tilde{X}; Y)$ imposes a partitioning problem on the model: it learns to split the entire probability space into regions with minimal overlap between different classes. As this problem is intractable for the large models considered in this work, we seek to find a local minima via stochastic gradient descent (SGD). We provide a further analysis of this loss function in Section 3.1.

2.3. Similarity Measure

Suppose that all data points in the training set are assigned their codewords according to the rule (3). Now we introduce how to compute a similarity between a query datapoint $\mathbf{x}^{(q)}$ and a support datapoint $\mathbf{x}^{(s)}$ for information retrieval or few-shot classification, where the superscripts $(q), (s)$ stand for query and support, respectively. Denote by $\tilde{\mathbf{x}}^{(s)}$ the codeword associated with $\mathbf{x}^{(s)}$, constructed by (3). For the test data $\mathbf{x}^{(q)}$, the encoder yields $p_{i,j}(\mathbf{x}^{(q)})$ for $i = 1, \dots, d$ and $j = 1, \dots, k$. As a similarity measure between $\mathbf{x}^{(q)}$ and $\mathbf{x}^{(s)}$, we calculate the following log probability.

$$\sum_{i=1}^d \log p_{i, \tilde{\mathbf{x}}_i^{(s)}}(\mathbf{x}^{(q)}). \quad (8)$$

The probabilistic quantity (8) indicates that $\mathbf{x}^{(q)}$ and $\mathbf{x}^{(s)}$ become more similar when the encoder's output—when $\mathbf{x}^{(q)}$ is provided—is well aligned with $\tilde{\mathbf{x}}^{(s)}$.

We can view our similarity measure (8) as a probabilistic generalization of the Hamming distance [8]. The Hamming distance quantifies the similarity between two strings of equal length as the number of positions at which the corresponding symbols are equal. As we have access to a distribution over codes, we use (8) to directly compute the log probability of having the same symbol at each position.

We use (8) as a similarity metric for both few-shot classification and image retrieval. We perform few-shot classification by computing a codeword for each class via (3) and classifying each test image by choosing the class that has the highest value of (8). We similarly perform image retrieval by mapping each support image to its most likely code (3) and for each query image retrieving the support image that has the highest (8).

While we have described the operations in (3) and (8) for a single pair $(\mathbf{x}^{(q)}, \mathbf{x}^{(s)})$, one can easily parallelize our evaluation procedure, since it is an argmax followed by a sum. Furthermore, $\tilde{\mathbf{x}}$ typically requires little memory, as it consists of discrete values, allowing us to compare against large support sets in parallel. Experiments in Section 5.4 investigate the degree of DIMCO's efficiency in terms of both time and memory.

2.4. Regularizing by Enforcing Independence

One way of interpreting the code distribution \tilde{X} is as a group of d separate code distributions $\tilde{x}_1, \dots, \tilde{x}_d$. Note that the similarity measure described in (8) can be seen as ensemble of the similarity measures of these d models. A classic result in ensemble learning is that using more diverse learners increases ensemble performance [9]. In a similar spirit, we used an optional regularizer which promotes pairwise independence between each pair in these d codes. Using this regularizer stabilized training, especially in more large-scale problems.

Specifically, we randomly sample pairs of indices i_1, i_2 from $\{1, \dots, d\}$ during each forward pass. Note that $\tilde{x}_{i_1} \otimes \tilde{x}_{i_2}$ and $(\tilde{x}_{i_1}, \tilde{x}_{i_2})$ are both categorical distributions with support size k^2 , and that we can estimate the two different distributions within each batch. We minimize their KL divergence to promote independence between these two distributions:

$$\mathcal{L}_{\text{ind}} = D_{\text{KL}}(\tilde{x}_{i_1} \otimes \tilde{x}_{i_2} \parallel \tilde{x}_{i_1}, \tilde{x}_{i_2}). \tag{9}$$

We compute (9) for a fixed number of random pairs of indices for each batch. The cost of computing this regularization term is miniscule compared to that of other components such as feeding data through the encoder.

Using this regularizer in conjunction with the learning objective (4) yields the following regularized loss:

$$\mathcal{L} = -I(\tilde{X}; Y) + \lambda \mathcal{L}_{\text{ind}} = -H(\tilde{X}) + H(\tilde{X}|Y) + \lambda \mathcal{L}_{\text{ind}}. \tag{10}$$

We fix $\lambda = 1$ in all experiments, as we found that DIMCO’s performance was not particularly sensitive to this hyperparameter. We emphasize that while this optional regularizer stabilizes training, our learning objective is the mutual information $I(\tilde{X}; Y)$ in (4).

2.5. Visualization of Codes

In Figure 2, we show images retrieved using our similarity measure (8). We trained a DIMCO model ($k = 16, d = 4$) on the CIFAR100 dataset. We selected specific code locations and plotted the top 10 test images according to our similarity measure. For example, the top (leftmost) image for code (\cdot, j_2, \cdot, j_4) would be computed as

$$\arg \max_{n \in \{1, \dots, N\}} \left[\log p_{2, j_2}(\mathbf{x}^n) + \log p_{4, j_4}(\mathbf{x}^n) \right], \tag{11}$$

where N is the number of test images.

We visualize two different combinations of codes in Figure 2. The two examples show that using codewords together results in their respective semantic concepts being combined: (man + fish = man holding fish), (round + warm color = orange). While we visualized combinations of 2 codewords for clarity, DIMCO itself uses a combination of d such codewords. The regularizer described in Section 2.4 further encourages each of these d codewords to represent different concepts. The combinatorially many (k^d) combinations in which DIMCO can assemble such codewords gives DIMCO sufficient expressive power to solve challenging tasks.

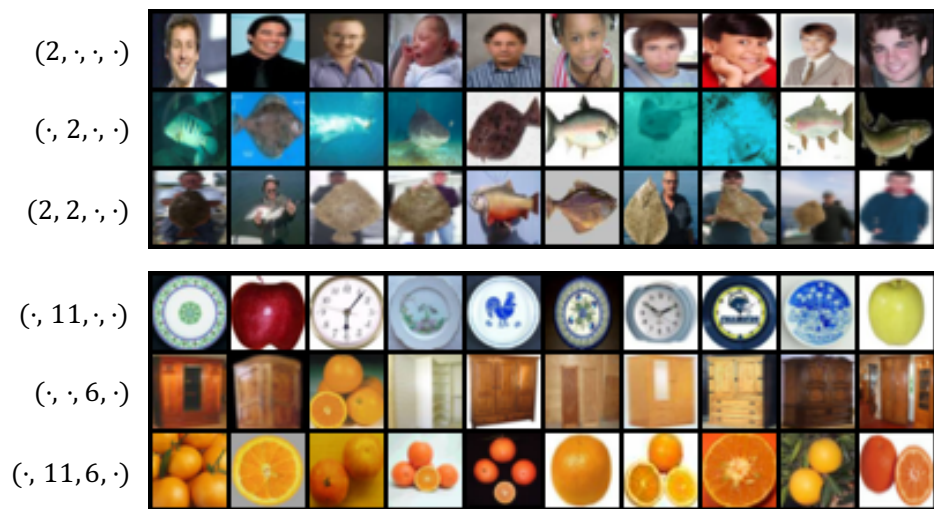


Figure 2. Compositionality of codes in a DIMCO model ($k = 16, d = 4$) trained on the CIFAR-100 dataset. Each row shows the top 10 images that assign highest marginal probability to specific codes (shown on left). Composed codewords retrieve images with combined semantic meaning, such as man + fish = man holding fish. We additionally visualize all kd codewords in Appendix A.

3. Analysis

3.1. Is Mutual Information a Good Objective?

Our learning objective for DIMCO (4) is the mutual information between codes and labels. In this subsection, we justify this choice by showing that many previous objectives are closely related to mutual information. Due to space constraints, we only show high-level connections here and provide a more detailed exposition in Appendix A.

3.1.1. Cross-Entropy

The de facto loss for classification is the cross-entropy loss, which is defined as

$$\text{xent}(Y, X) = -\mathbb{E}_{y \sim Y, x \sim X} \left[\log q(\hat{Y} = y | \tilde{X}(x)) \right], \tag{12}$$

where \hat{Y} is the model’s prediction of Y . Using the observation that the final layer $q(\cdot)$ acts a parameterized approximation of the true conditional distribution $p(Y|\tilde{X})$, we write this as

$$\begin{aligned} \text{xent}(Y, X) &\approx \mathbb{E}_{y \sim Y, x \sim X} [-\log p(\mathbf{y}|\tilde{\mathbf{x}})] \\ &= -I(\tilde{X}; Y) + H(Y). \end{aligned} \tag{13}$$

The $H(Y)$ term can be ignored since it is not affected by model parameters. Therefore, minimizing cross-entropy is approximately equivalent to maximizing mutual information. The two objectives become completely equivalent when the final linear layer $q(\cdot)$ perfectly represents the conditional distribution $q(\mathbf{y}|\tilde{\mathbf{x}})$. Note that for discrete $\tilde{\mathbf{x}}$, we cannot use a linear layer to parameterize $q(\mathbf{y}|\tilde{\mathbf{x}})$, and therefore, cannot directly optimize the cross-entropy loss. We can therefore view our loss as a necessary modification of the cross-entropy loss for our setup of using discrete embeddings.

3.1.2. Contrastive Losses

Many metric learning methods [1,2,10–12] use a contrastive learning objective to learn a continuous embedding (\tilde{X}). Such contrastive losses consist of (1) a positive term that encourages an embedding to move closer to that of other relevant embeddings and (2) a negative term that encourages it to move away from irrelevant embeddings. The

positive term approximately minimizes $\log p(\tilde{X}|y)$, and the negative term as approximately minimizes $\log p(\tilde{X})$. Together, these terms have the combined effect of maximizing

$$\begin{aligned} &\mathbb{E}[\log p(\tilde{x}|y) - \log p(\tilde{x})] \\ &= -H(\tilde{X}|Y) + H(\tilde{X}) = I(\tilde{X}; Y). \end{aligned} \tag{14}$$

We show such equivalences in detail in Appendix A.

In addition to these direct connections to previous loss functions, we show empirically in Section 5.1 that the mutual information strongly correlates with both the top-1 accuracy metric for classification and the Recall@1 metric for retrieval.

3.2. Does Using Discrete Codes Help Generalization?

In Section 1, we have provided motivation for the use of discrete codes through the regularization effect of an information bottleneck. In this subsection, we theoretically analyze whether learning discrete codes by maximizing mutual information leads to better generalization. In particular, we study how the mutual information on the test set is affected by the choice of input dataset structure and code hyperparameters k and d through a PAC learning bound.

We analyze DIMCO’s characteristics at the level of minibatches. Following related meta-learning works [13,14], we call each batch a “task”. We note that this is only a difference in naming convention, and our analysis applies equally well to the metric learning setup: we can view each batch consisting of support and query points as a task.

Define a task T to be a distribution over $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. Let tasks T^1, \dots, T^n be sampled i.i.d. from a distribution of tasks τ . Each task T consists of a fixed-size dataset $D_T = z_T^1, \dots, z_T^m = (x_T^1, y_T^1), \dots, (x_T^m, y_T^m)$, which is a set of m i.i.d. samples from the data distribution ($z_T^j \sim T$). Let θ be the parameters of DIMCO. Let X, Y, \tilde{X} be the random variables for data, labels, and codes, respectively. Recall that our objective is the expected mutual information between labels and codes:

$$\mathcal{L}(\tau, \theta) = -\mathbb{E}_{T \sim \tau} [I(\tilde{X}(X_T, \theta); Y_T)]. \tag{15}$$

The loss that we actually optimize (Equations (6) and (7)) is the empirical loss:

$$\hat{\mathcal{L}}(T^{1:n}, \theta) = -\frac{1}{n} \sum_{i=1}^n \hat{I}(\tilde{X}(X_{T^i}, \theta); Y_{T^i}). \tag{16}$$

The following theorem bounds the difference between the expected loss \mathcal{L} and the empirical loss $\hat{\mathcal{L}}$.

Theorem 1. *Let d_Θ be the VC dimension of the encoder $\tilde{X}(\cdot)$. The following inequality holds with high probability:*

$$\mathcal{L}(\tau, \theta) - \hat{\mathcal{L}}(T^{1:n}, \theta) \leq O\left(\sqrt{\frac{d_\Theta}{n} \log \frac{n}{d_\Theta}}\right) + O\left(\frac{|\tilde{X}| \log(m)}{\sqrt{m}}\right) + O\left(\frac{|\tilde{X}||Y|}{m}\right). \tag{17}$$

Proof. We use VC dimension bounds and a finite sample bound for mutual information [15]. We defer a detailed statement and proof to Appendix B. \square

First note that all three terms in our generalization gap (A11) converge to zero as $n, m \rightarrow \infty$. This shows that training a model by maximizing empirical mutual information, as in Equations (6) and (7), generalizes perfectly in the limit of infinite data.

Theorem 1 also shows how the generalization gap is affected differently by dataset size m and number of datasets n . A large n directly compensates for using a large backbone (d_Θ), and a large m compensates for using a large final representation (\tilde{X}). Put differently,

to effectively learn from small datasets (m), one should use a small representation (\tilde{X}). The number of datasets n is typically less of a problem because the number of different ways to sample datasets is combinatorially large (e.g., $n > 10^{10}$ for miniImagenet 5-way 1-shot tasks). Recall that DIMCO has $|\tilde{X}| = d \log k$, meaning that we can control the latter two terms using our hyperparameters d, k . We have explained the use of discrete codes through the information bottleneck effect of small codes \tilde{X} , and Theorem 1 confirms this intuition.

4. Related Work

Information bottleneck. DIMCO and Theorem 1 are both close in spirit to the information bottleneck (IB) principle [5,6,16]. IB finds a set of compact representatives \tilde{X} while maintaining sufficient information about Y , minimizing the following objective function:

$$\mathcal{J}(p(\tilde{\mathbf{x}}|\mathbf{x})) = I(\tilde{X}; X) - \beta I(\tilde{X}; Y), \quad (18)$$

subject to $\sum_{\tilde{\mathbf{x}}} p(\tilde{\mathbf{x}}|\mathbf{x}) = 1$. Equivalently, it can be stated that one maximizes $I(\tilde{X}; Y)$ while simultaneously minimizing $I(\tilde{X}; X)$. Similarly, our objective (15) is information maximization $I(\tilde{X}; Y)$, and our bound (A11) suggests that the representation capacity $|\tilde{X}|$ should be low for generalization. In the deterministic information bottleneck [17], $I(\tilde{X}; X)$ is replaced by $H(\tilde{X})$. These three approaches to generalization are related via the chain of inequalities $I(\tilde{X}; X) \leq H(\tilde{X}) \leq \log |\tilde{X}|$, which is tight in the limit of \tilde{X} being incompressible. For any finite representation, i.e., $|\tilde{X}| = N$, the limit $\beta \rightarrow \infty$ in (18) yields a hard partitioning of X into N disjoint sets. DIMCO uses the infomax principle to learn $N = k^d$ such representatives, which are arranged by k -way d -dimensional discrete codes for compact representation with sufficient information on Y .

Regularizing meta-learning. Previous meta-learning methods have restricted task-specific learning by learning only a subset of the network [18], learning on a low-dimensional latent space [19], learning on a meta-learned prior distribution of parameters [20], and learning context vectors instead of model parameters [21]. Our analysis in Theorem 1 suggests that reducing the expressive power of the task-specific learner has a meta-regularizing effect, indirectly giving theoretical support for previous works that benefited from reducing the expressive power of task-specific learners.

Discrete representations. Discrete representations have been thoroughly studied in information theory [22]. Recent deep learning methods directly learn discrete representations by learning generative models with discrete latent variables [23–25] or maximizing the mutual information between representation and data [26]. DIMCO is related to but differs from these works, as it assumes a supervised meta-learning setting and performs infomax using *labels* instead of data.

A standard approach to learning label-aware discrete codes is to first learn continuous embeddings and then quantize it using an objective that maximally preserves its information [27–29]. DIMCO can be seen as an end-to-end alternative to quantization which directly learns discrete codes. Jeong and Song [30] similarly learns a sparse binary code in an end-to-end fashion by solving a minimum cost flow problem with respect to labels. Their method differs from DIMCO, which learns a dense discrete code by optimizing $I(\tilde{X}; Y)$, which we estimate with a closed-form formula.

Metric learning. The structure and loss function of DIMCO are closely related to those of metric learning methods [1,11,12,31]. We show that the loss functions of these methods can be seen as approximations of the mutual information ($I(\tilde{X}; Y)$) in Section 2.2, and provide more in-depth exposition in Appendix A. While all of these previous methods require a support/query split within each batch, DIMCO simply optimizes an information-theoretic quantity of each batch, removing the need for such structured batch construction.

Information theory and representation learning. Many works have applied information-theoretic principles to unsupervised representation learning: to derive an objective for GANs to learn disentangled features [32], to analyze the evidence lower bound (ELBO) [33,34], and to directly learn representations [35–39]. Related also are previous methods that enforce independence within an embedding [40,41]. DIMCO is also an information-theoretic representation learning method, but we instead assume a supervised learning setup where the representation must reflect ground-truth labels. We also used previous results from information theory to prove a generalization bound for our representation learning method.

5. Experiments

In our experiments, we used datasets with varying degrees of complexity: CIFAR10/100 [42], miniImageNet [31], CUB200 [43], Cars196 [44], and ImageNet (ILSVRC-2012-CLS, Deng et al. [45]). We used standard train/test splits for each dataset unless stated otherwise. We also used various network architectures: 4-layer convnet [31] and ResNet12/20/50 [46,47]. We followed previously reported experimental setups as closely as possible, and provide minor experiment details in Appendix C.

5.1. Correlation of Metrics

We have shown in Section 3.1 that the mutual information $I(\tilde{X}; Y)$ is strongly connected to previous loss functions for classification and retrieval. In this subsection, we show experiments performed to verify whether $I(\tilde{X}; Y)$ is a good *metric* that quantitatively shows the quality of the representation \tilde{X} . We trained DIMCO on the miniImageNet dataset with $k = d = 64$ for 20 epochs. We plot the pairwise correlations between five different metrics: (5, 10, 20)-way 1-shot accuracy, Recall@1, and $I(\tilde{X}; Y)$. The results in Figure 3 show that all five metrics are very strongly correlated. We observed similar trends when training with loss functions other than $I(\tilde{X}; Y)$ as well; we show these experiments in Appendix C due to space constraints.

5.2. Label-Aware Compression

We applied DIMCO to compressing feature vectors of trained classifier networks. We obtained penultimate embeddings of ResNet20 networks each trained on CIFAR10 and CIFAR100. The two networks had top-1 accuracies of 91.65 and 66.61, respectively. We trained on embeddings for the train set of each dataset, and measured top-1 accuracy of the test set using the training set as support. We compare DIMCO to product quantization (PQ, Jégou et al. [28]), which similarly compresses a given embedding to a k -way d -dimensional code. We compare the two methods in Table 1 with the same range of k, d hyperparameters. We performed the same experiment on the larger ImageNet dataset with a ResNet50 network which had a top-1 accuracy of 76.00. We compare DIMCO to both adaptive scalar quantization (SQ) and PQ in Table 2. We show extended experiments for all three datasets in Appendix A.

The results in Tables 1 and 2 demonstrate that DIMCO consistently outperforms PQ, and is especially efficient when d is low. Furthermore, the ImageNet experiment (Table 2) shows that DIMCO even outperforms SQ, which has a much lower compression rate compared to the embedding sizes we consider for DIMCO. These results are likely due to DIMCO performing *label-aware compression*, where it compresses the embedding while taking the label into account, whereas PQ and SQ only compress the embeddings themselves.

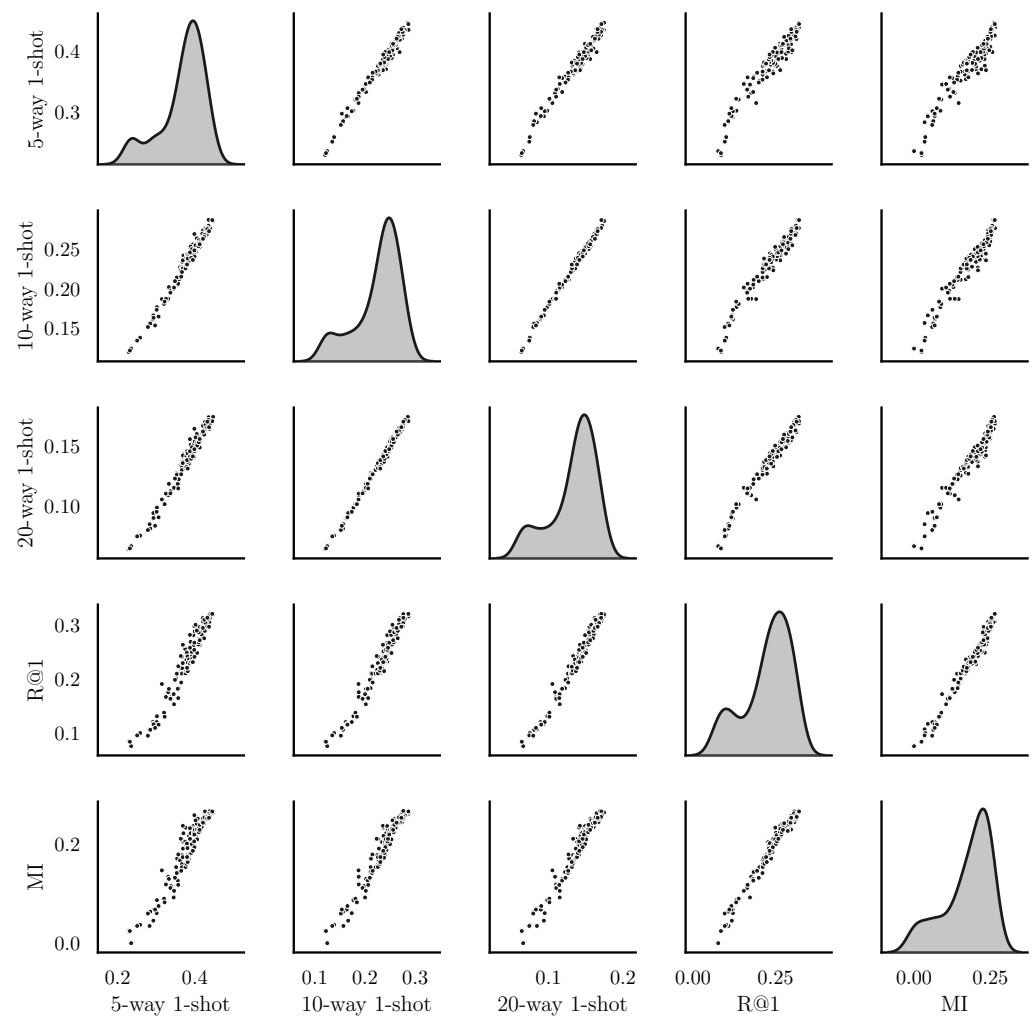


Figure 3. Pairwise correlations between $MI = I(X; \tilde{X})$ and previous metrics. Best viewed zoomed in.

Table 1. Top-1 accuracies under various compressed CIFAR-10/100 embeddings with size (k, d) , with best results for each setting in bold. The original embeddings were continuous 64-dimensional embeddings ($k \approx 2^{32}, d = 64$).

		CIFAR-10		CIFAR-100	
		$d = 4$	$d = 16$	$d = 4$	$d = 16$
PQ	$k = 2$	50.38	87.24	7.77	32.47
	$k = 4$	86.86	90.92	18.78	52.02
	$k = 8$	90.44	91.49	31.04	58.82
	$k = 16$	91.02	91.52	43.16	61.00
DIMCO (ours)	$k = 2$	64.43	88.85	10.03	33.46
	$k = 4$	88.47	91.04	25.7	53.36
	$k = 8$	91.29	91.45	46.22	58.84
	$k = 16$	91.68	91.57	57.83	62.49

Table 2. Top-1 accuracies and compression rates of ImageNet embeddings under various compressed embedding sizes (k, d), with best results in bold. The compression rate is the ratio between uncompressed and compressed sizes; it is calculated as $\frac{32 \cdot 2048}{d \log k}$.

Method	(k, d)	Compression Rate	Accuracy
SQ	(2, 2048)	32	0.10
SQ	(4, 2048)	16	0.28
SQ	(8, 2048)	10	12.36
SQ	(16, 2048)	8	57.80
PQ	(2, 2)	32,768	0.21
PQ	(4, 4)	8192	0.92
PQ	(8, 8)	2730	11.91
PQ	(16, 16)	1024	44.93
DIMCO	(2, 2)	32,768	0.14
DIMCO	(4, 4)	8192	2.68
DIMCO	(8, 8)	2730	33.28
DIMCO	(16, 16)	1024	63.64

5.3. Few-Shot Classification

We evaluated DIMCO’s few-shot classification performance on the miniImageNet dataset. We compare our method against the following previous works: Snell et al. [3], Vinyals et al. [31], Liu et al. [48], Ye et al. [49], Ravi and Larochelle [50], Sung et al. [51], Bertinetto et al. [52], Lee et al. [53]. All methods use the standard four-layer convnet with 64 filters per layer. While some methods used more filters, we used 64 for a fair comparison. We used the data augmentation scheme proposed by Lee et al. [53] and used balanced batches of 100 images consisting of 10 different classes. We evaluated both 5-way 1-shot and 5-way 5-shot learning, and report 95% confidence intervals of 1000 random episodes on the test split.

Results are shown in Table 3, and we provide an extended table with an alternative backbone in Appendix A. Figure 3 shows that DIMCO outperforms previous works on the 5-way 1-shot benchmark. DIMCO’s 5-way 5-shot performance is relatively low, likely because the similarity metric (Section 2.3) handles support datapoints individually instead of aggregating them, similarly to Matching Nets [31]. Additionally, other methods are explicitly trained to optimize 5-shot performance, whereas DIMCO’s training procedure is the same regardless of task structure.

Table 3. Few-shot classification accuracies on the miniImageNet benchmark, with best results for each setting in bold. † denotes transductive methods, which are more expressive by taking unlabeled examples into account.

Method	5-Way 1-Shot	5-Way 5-Shot
† TPN	55.51 ± 0.86	69.86 ± 0.65
† FEAT	55.75 ± 0.20	72.17 ± 0.16
MetaLSTM	43.44 ± 0.77	60.60 ± 0.71
MatchingNet	43.56 ± 0.84	55.31 ± 0.73
ProtoNet	49.42 ± 0.78	68.20 ± 0.66
RelationNet	50.44 ± 0.82	65.32 ± 0.70
R2D2	51.2 ± 0.6	68.8 ± 0.1
MetaOptNet-SVM	52.87 ± 0.57	68.76 ± 0.48
DIMCO (64, 64)	47.33 ± 0.46	61.59 ± 0.52
DIMCO (64, 128)	53.29 ± 0.47	64.79 ± 0.57

5.4. Image Retrieval

We conducted image retrieval experiments using two standard benchmark datasets: CUB-200-2011 and Cars-196. As baselines, we used three widely adopted metric learning methods: Binomial Deviance [54], Triplet loss [1], and Proxy-NCA [2]. The backbone for all methods was a ResNet-50 network pretrained on the ImageNet dataset. We trained DIMCO on various combinations of (p, d) , and set the embedding dimension of the baseline methods to 128. We measured the time per query for each method on a Xeon E5-2650 CPU without any parallelization. We note that computing the retrieval time using a parallel implementation would skew the results even more in favor of DIMCO, since DIMCO’s evaluation is simply one memory access followed by a sum.

Results presented in Table 4 show that DIMCO outperforms all three baseline, and that the compact code of DIMCO takes roughly an order of magnitude less memory, and requires less query time as well. This experiment also demonstrates that discrete representations can outperform modern methods that use continuous embeddings, even on this relatively large-scale task. Additionally, this experiment shows that DIMCO can train using large backbones without significantly overfitting.

Table 4. Image retrieval performance on CUB-200-2011 and Cars-196, measured by Recall@1, with best results for each setting in bold. Memory is the number of bits that an embedding vector of each image uses. Time is seconds taken to retrieve a single query from database (5924 and 8131 images for CUB-200-2011 and Cars-196, respectively).

Method	(k, d)	Memory [Bits]	CUB-200-2011		Cars-196	
			Recall@1	Time [s]	Recall@1	Time [s]
Binomial Deviance	(-, 128)	4096	57.25	16.37	72.53	21.86
Triplet	(-, 128)	4096	56.80	16.37	73.79	21.86
Proxy-NCA	(-, 128)	4096	56.19	16.37	75.94	21.86
DIMCO (ours)	(32, 32)	160	51.04	1.48	63.44	2.64
	(64, 64)	384	55.78	2.85	72.06	6.01
	(128, 128)	896	58.05	5.81	76.04	11.92
	(256, 256)	2048	58.90	12.20	77.32	16.04

6. Discussion

We introduced DIMCO, a model that learns a discrete representation of data by directly optimizing the mutual information with the label. To evaluate our initial intuition that shorter representations generalize better between tasks, we provided generalization bounds that get tighter as the representation gets shorter. Our experiments demonstrated that DIMCO is effective at both compressing a continuous embedding, and also at learning a discrete embedding from scratch in an end-to-end manner. The discrete embeddings of DIMCO outperformed recent continuous feature extraction methods while also being more efficient in terms of both memory and time. We believe the tradeoff between discrete and continuous embeddings is an exciting area for future research.

DIMCO was motivated by concepts such as the minimum description length (MDL) principle and the information bottleneck: compact task representations should have less room to overfit. Interestingly, Yin et al. [55] reports that doing the opposite—regularizing the task-general parameters—prevents meta-overfitting by discouraging the meta-learning model from memorizing the given set of tasks. In future work, we will investigate the common principle underlying these seemingly contradictory approaches for a fuller understanding of meta-generalization.

Author Contributions: Conceptualization, Y.L., W.K. and S.C.; methodology, Y.L., W.K. and S.C.; software, Y.L., W.K. and W.P.; validation, Y.L., W.K., W.P. and S.C.; formal analysis, Y.L. and S.C.; investigation, Y.L., W.K., W.P. and S.C.; resources, Y.L.; writing—original draft preparation, Y.L.; writing—review and editing, Y.L., W.K., W.P. and S.C.; visualization, Y.L., W.K. and W.P.; supervision, S.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Previous Loss Functions Are Approximations of Mutual Information

Appendix A.1. Cross-Entropy Loss

The cross-entropy loss has directly been used for few-shot classification [3,31].

Let $q(\mathbf{y}|\tilde{\mathbf{x}}; \phi)$ be a parameterized prediction of \mathbf{y} given $\tilde{\mathbf{x}}$, which tries to approximate the true conditional distribution $q(\mathbf{y}|\tilde{\mathbf{x}})$. Typically, in a classification network, ϕ is the parameters of a learned projection matrix and $q(\cdot)$ is the final linear layer. The expected cross-entropy loss can be written as

$$\text{xent}(Y, \tilde{X}) = \mathbb{E}_{\mathbf{y} \sim Y, \tilde{\mathbf{x}} \sim \tilde{X}} [-\log q(\mathbf{y}|\tilde{\mathbf{x}}, \phi)]. \tag{A1}$$

Assuming that the approximate distribution $q(\cdot)$ is sufficiently close to $p(\mathbf{y}|\tilde{\mathbf{x}})$, minimizing (A1) can be seen as

$$\arg \min \text{xent}(Y, \tilde{X}) \approx \arg \min \mathbb{E}_{\mathbf{y} \sim Y, \tilde{\mathbf{x}} \sim \tilde{X}} [-\log p(\mathbf{y}|\tilde{\mathbf{x}})] \tag{A2}$$

$$= \arg \min H(Y|\tilde{X}) = \arg \max I(\tilde{X}; Y), \tag{A3}$$

where the last equality uses the fact that $H(Y)$ is independent of model parameters. Therefore, cross-entropy minimization is approximate maximization of the mutual information between representation \tilde{X} and labels Y .

The approximation is that we parameterized $q(\mathbf{y}|\tilde{\mathbf{x}}; \phi)$ as a linear projection. This structure cannot generalize to new classes because the parameters ϕ are specific to the labels \mathbf{y} seen during training. For a model to generalize to unseen classes, one must amortize the learning of this approximate conditional distribution. [3,31] sidestepped this issue by using the embeddings for each class as ϕ .

Appendix A.2. Triplet Loss

The triplet loss [1] is defined as

$$\mathcal{L}_{\text{triplet}} = \left\| \tilde{\mathbf{x}}_q - \tilde{\mathbf{x}}_p \right\|_2^2 - \left\| \tilde{\mathbf{x}}_q - \tilde{\mathbf{x}}_n \right\|_2^2, \tag{A4}$$

where $\tilde{\mathbf{x}}_q, \tilde{\mathbf{x}}_p, \tilde{\mathbf{x}}_n \in \mathbb{R}^d$ are the embedding vectors of query, positive, and negative images. Let \mathbf{y}_q denote the label of the query data. Recall that the pdf function of a unit Gaussian is $\log N(\tilde{\mathbf{x}}|\mu, 1) = -c_1 - c_2 \|\tilde{\mathbf{x}} - \mu\|_2^2$, where c_1, c_2 are constants. Let $p_p(\tilde{\mathbf{x}}) = N(\tilde{\mathbf{x}}_p, 1)$ and $p_n(\tilde{\mathbf{x}}) = N(\tilde{\mathbf{x}}_n, 1)$ be unit Gaussian distributions centered at $\tilde{\mathbf{x}}_p, \tilde{\mathbf{x}}_n$, respectively. We have

$$\mathbb{E}[-\mathcal{L}_{\text{triplet}}] \propto \mathbb{E}[\log p_p(\tilde{\mathbf{x}}) - \log p_n(\tilde{\mathbf{x}})] \tag{A5}$$

$$\approx \mathbb{E}[\log p_p(\tilde{\mathbf{x}}) - \log p(\tilde{\mathbf{x}})] \tag{A6}$$

$$= -H(\tilde{X}|Y) + H(\tilde{X}) = I(\tilde{X}; Y). \tag{A7}$$

Two approximations were made in the process. We first assumed that the embedding distribution of images not in \mathbf{y}_q is equal to the distribution of all embeddings. This is reasonable when each class only represents a small fraction of the full data. We also

approximated the embedding distributions $p(\tilde{\mathbf{x}}|\mathbf{y}), p(\tilde{\mathbf{x}})$ with unit Gaussian distributions centered at single samples from each.

Appendix A.3. N-Pair Loss

Multiclass N -pair loss [11] was proposed as an alternative to Triplet loss. This loss function requires one positive embedding $\tilde{\mathbf{x}}^+$ and multiple negative embeddings $\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_{N-1}$, and takes the form

$$-\log \frac{\exp(\tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}^+)}{\exp(\tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}^+) + \sum_{i=1}^{N-1} \exp(\tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}_i)}. \tag{A8}$$

This can be seen as the cross-entropy loss applied to $\text{softmax}(\tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}^+, \tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}_{N-1})$.

Following the same logic as the cross-entropy loss, this is also an approximation of $I(\tilde{X}; Y)$. This objective should have less variance than Triplet loss since it approximates $p(\tilde{\mathbf{x}})$ using more examples.

Appendix A.4. Adversarial Metric Learning

Deep Adversarial Metric Learning [12] tackles the problem of most negative examples being uninformative by directly generating meaningful negative embeddings. This model employs a *generator* which takes as input the embeddings of anchor, positive, and negative images. The generator then outputs a “synthetic negative” embedding that is hard to distinguish from a positive embedding while being close to the negative embedding.

This can be seen as optimizing

$$\mathbb{E} \left[\log p_p(\tilde{\mathbf{x}}) - \log p(\tilde{\mathbf{x}}) \right] = I(\tilde{X}; Y) \tag{A9}$$

by estimating $p(\tilde{\mathbf{x}})$ using a generative network rather than directly from samples. Rather than modelling the marginal distribution $p(\tilde{\mathbf{x}})$, this method conditionally models $p(\tilde{\mathbf{x}}; \tilde{\mathbf{x}}_q, \tilde{\mathbf{x}}_p, \tilde{\mathbf{x}}_n)$ so that $\tilde{\mathbf{x}}$ is hard to distinguish from $\tilde{\mathbf{x}}_p$ while sufficiently close to both $\tilde{\mathbf{x}}_q$ and $\tilde{\mathbf{x}}_n$.

Appendix B. Proof of Theorem 1

We restate and prove our main theorem.

Theorem A1. *Let d_Θ be the VC dimension of the encoder $\tilde{X}(\cdot)$. Let $\hat{I}(\tilde{X}(X_T, \theta); Y_T)$ be the empirical estimate of the mutual information using finite dataset D_T , and define empirical loss as*

$$\hat{\mathcal{L}}(T^{1:n}, \theta) = -\frac{1}{n} \sum_{i=1}^n \hat{I}(\tilde{X}(X_{T^i}, \theta); Y_{T^i}). \tag{A10}$$

The following inequality holds with high probability:

$$\mathcal{L}(\tau, \theta) - \hat{\mathcal{L}}(T^{1:n}, \theta) \leq O\left(\sqrt{\frac{d_\Theta}{n} \log \frac{n}{d_\Theta}}\right) + O\left(\frac{|\tilde{X}| \log(m)}{\sqrt{m}}\right) + O\left(\frac{|\tilde{X}||Y|}{m}\right) \tag{A11}$$

Proof. We use the following lemma from [15], which we restate using our notation.

Lemma A1. *Let \tilde{X} be a random mapping of X . Let D be a sample of size m drawn from the joint probability distribution $p(X, Y)$. Denote the empirical mutual information observed from D between \tilde{X} and Y as $\hat{I}(\tilde{X}; Y)$. For any $\delta \in (0, 1)$, the following holds with probability at least $1 - \delta$:*

$$|I(\tilde{X}; Y) - \hat{I}(\tilde{X}; Y)| \leq \frac{(3|\tilde{X}| + 2) \log(m) \sqrt{\log(4/\delta)}}{\sqrt{2m}} + \frac{(|Y| + 1)(|\tilde{X}| + 1) - 4}{m} \tag{A12}$$

We simplify this and plug in our specific quantities of interest $(\tilde{X}(X_T, \theta), Y_T)$:

$$\left| I(\tilde{X}(X_T, \theta); Y_T) - \hat{I}(\tilde{X}(X_T, \theta); Y_T) \right| \leq O\left(\frac{|\tilde{X}| \log(m)}{\sqrt{m}}\right) + O\left(\frac{|\tilde{X}||Y|}{m}\right). \tag{A13}$$

We similarly bound the error caused by estimating \mathcal{L} with a finite number of tasks sampled from τ . Denote the finite sample estimate of \mathcal{L} as

$$\hat{\mathcal{L}}(\tau, \theta) = -\frac{1}{n} \sum_{i=1}^n I(\tilde{X}(X_{T_i}, \theta); Y_{T_i}). \tag{A14}$$

Let the mapping $X \mapsto \tilde{X}$ be parameterized by $\theta \in \Theta$ and let this model have VC dimension d_Θ . Using d_Θ , we can state that with high probability,

$$\left| \mathcal{L}(\tau, \theta) - \hat{\mathcal{L}}(\tau, \theta) \right| \leq O\left(\sqrt{\frac{d_\Theta}{n} \log \frac{n}{d_\Theta}}\right), \tag{A15}$$

where d_Θ is the VC dimension of hypothesis class Θ .

Combining equations (A15, A13), we have with high probability

$$\left| \mathcal{L}(\tau, \theta) - \left(-\frac{1}{n} \sum_{i=1}^n \hat{I}(\tilde{X}(X_{T_i}, \theta); Y_{T_i})\right) \right| \tag{A16}$$

$$\leq \left| \mathcal{L}(\tau, \theta) - \hat{\mathcal{L}}(\tau, \theta) \right| + O\left(\frac{|\tilde{X}| \log(m)}{\sqrt{m}}\right) + O\left(\frac{|\tilde{X}||Y|}{m}\right) \tag{A17}$$

$$\leq O\left(\sqrt{\frac{d_\Theta}{n} \log \frac{n}{d_\Theta}}\right) + O\left(\frac{|\tilde{X}| \log(m)}{\sqrt{m}}\right) + O\left(\frac{|\tilde{X}||Y|}{m}\right). \tag{A18}$$

□

Appendix C. Experiments

Appendix C.1. Parameterizing the Code Layer

Recall that each discrete code is parameterized by a $\mathbb{R}^{k \times d}$ matrix. A problem with a naive implementation of DIMCO is that simply using a linear layer that maps \mathbb{R}^D to $\mathbb{R}^{k \times d}$ takes Dkd parameters in that single layer. This can be prohibitively expensive for large embeddings, e.g., $d = 4096$. We therefore, parameterize this code layer as the product of two matrices, which sequentially map $\mathbb{R}^D \rightarrow \mathbb{R}^n \rightarrow \mathbb{R}^{k \times d}$. The total number of parameters required for this is $nD + nkd$. We fix all $n = 128$. While more complicated tricks could reduce the parameter count even further, we found that this simple structure was sufficient to produce the results in this paper.

Appendix C.2. Correlation of Metrics

We collected statistics from 8 different independent runs, and report the averages of 500 batches of 1-shot accuracies, Recall@1, and mutual information. $I(\tilde{X}; Y)$ was computed using balanced batches of 16 images each from 5 different classes. In addition to the experiment in the paper, we measured the correlations between 1-shot accuracies, Recall@1, and NMI using three previously proposed losses (triplet, npair, protonet). Figure A1 shows that even for other methods for which mutual information is not the objective, mutual information strongly correlates with all other previous metrics.

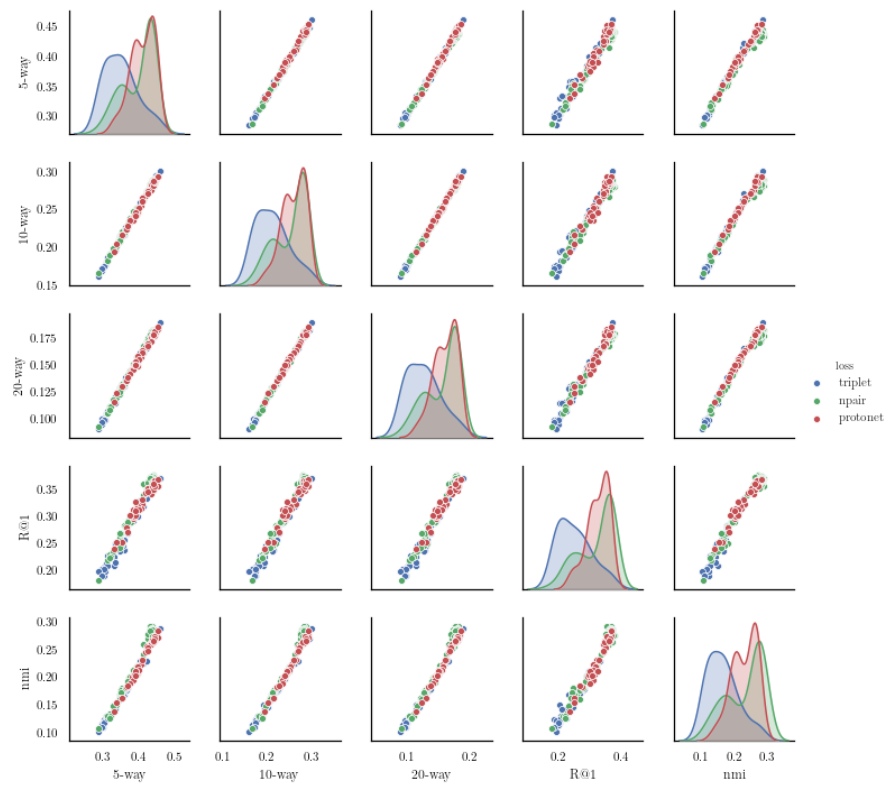


Figure A1. Correlation between few-shot accuracy and retrieval metrics.

Appendix C.3. Code Visualizations

We provide additional visualizations of codes in Figure A2. These examples consistently show that each code encodes a semantic concept, and that such concepts can be but are not necessarily tied to a particular class.

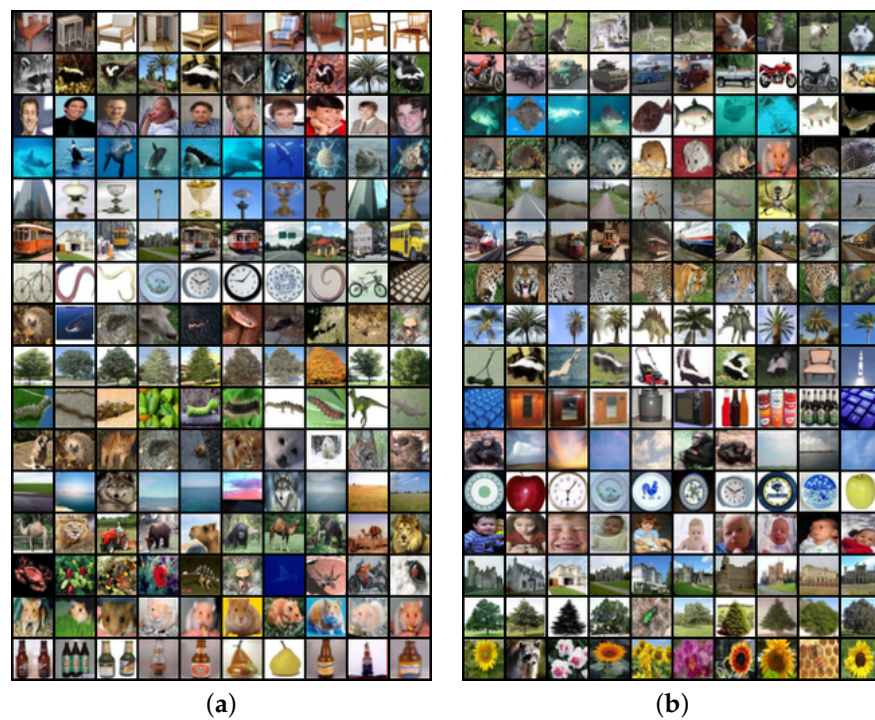


Figure A2. Cont.

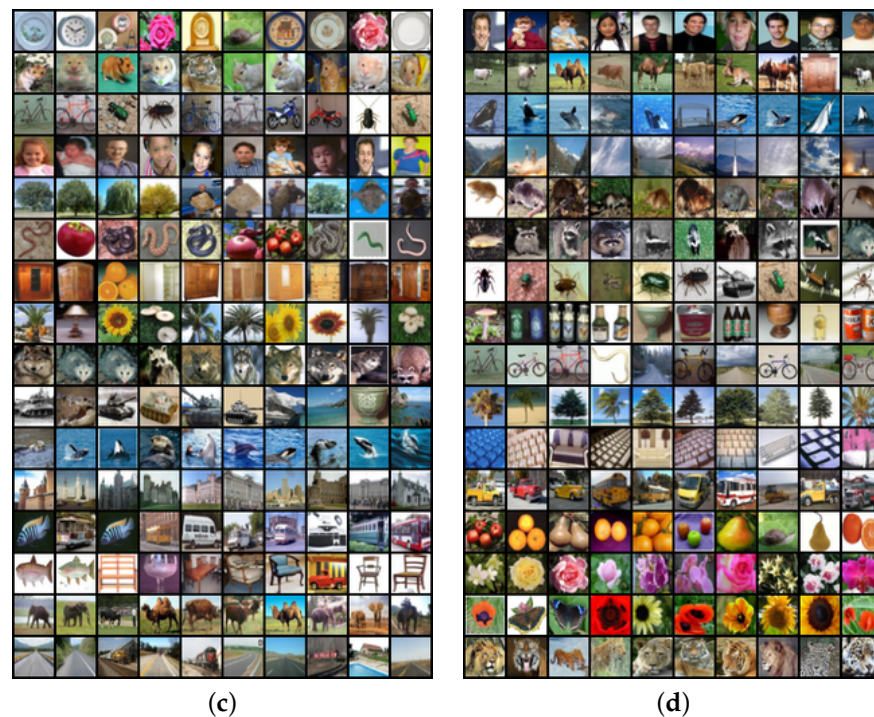


Figure A2. Visualization of codes of a DIMCO model ($k = 16, d = 4$) trained on CIFAR100. Each of the subfigures (a–d) correspond to one of the $d = 4$ dimensions. For each of the $kd = 64$ codewords, we show the top 10 images from the *test set* that assign it the highest marginal probability in each row.

Table A1. Top-1 accuracies under various compressed embedding sizes (k, d), with best results for each setting in bold. The original embeddings were continuous 64-dimensional embeddings, which corresponds to $k = 2^{32}, d = 64$.

		CIFAR-10				CIFAR-100			
		$d = 2$	$d = 4$	$d = 8$	$d = 16$	$d = 2$	$d = 4$	$d = 8$	$d = 16$
Product Quantization	$k = 2$	21.88	50.38	78.68	87.24	2.84	7.77	15.24	32.47
	$k = 4$	60.70	86.86	89.53	90.92	7.69	18.78	35.62	52.02
	$k = 8$	90.46	90.44	90.98	91.49	15.28	31.04	50.74	58.82
	$k = 16$	91.29	91.02	91.27	91.52	25.71	43.16	55.88	61.00
DIMCO (ours)	$k = 2$	36.37	64.43	83.10	88.85	3.66	10.03	19.31	33.46
	$k = 4$	62.92	88.47	90.78	91.04	11.83	25.7	37.4	53.36
	$k = 8$	90.77	91.29	91.41	91.45	31.9	46.22	52.77	58.84
	$k = 16$	91.49	91.68	91.46	91.57	46.17	57.83	61.11	62.49

Appendix C.4. Label-Aware Bit Compression

We computed top-1 accuracies using a kNN classifier on each type of embedding with $k = 200$. We present extended results comparing PQ and DIMCO on ImageNet embeddings in Table A2. For CIFAR-10 and CIFAR-100 pretrained ResNet20, we used pretrained weights of open-sourced repository (<https://github.com/cheneafo/CIFAR-pretrained-models>, accessed on 15 March 2021), and for ImageNet pretrained ResNet50, we used torchvision. We optimized the probabilistic encoder with Adam optimizer [56] with learning rate of 1×10^{-2} for CIFAR-100 and ImageNet, and 3×10^{-3} for CIFAR-10.

Table A2. Top-1 accuracies of compressed embeddings under various sizes (k, d) on the ImageNet dataset. The original embeddings were continuous 2048-dimensional embeddings, which corresponds to $k = 2^{32}, d = 2048$.

	Product Quantization				DIMCO			
	$d = 2$	$d = 4$	$d = 8$	$d = 16$	$d = 2$	$d = 4$	$d = 8$	$d = 16$
$k = 2$	0.21	0.31	1.00	5.21	0.14	0.42	1.47	5.50
$k = 4$	0.20	0.92	4.86	21.45	0.67	2.68	11.80	26.96
$k = 8$	0.82	2.57	11.91	35.43	1.71	12.82	33.28	48.11
$k = 16$	1.99	5.66	20.34	44.93	4.90	26.22	44.01	55.34

Appendix C.5. Small Train Set

We performed an experiment to see how well DIMCO can generalize to new datasets *after training with a small number of datasets*. We trained each model using $\{1, 4, 16, 64\}$ samples from each class in the miniImageNet dataset. For example, 4 samples means that we trained on (64 classes \times 4 images) instead of the full (64 classes \times 600 images). We compare our method against three methods which use continuous embeddings for each datapoint: Triplet Nets [1], multiclass N-pair loss [11], and ProtoNets [3].

Figure A3 shows that DIMCO learns much more effectively compared to previous methods when the number of examples per class is low. We attribute this to DIMCO’s tight generalization gap. Since DIMCO uses fewer bits to describe each datapoint, the codes act as an implicit regularizer that helps generalization to unseen datasets. We additionally note that DIMCO is the only method in Figure A3 that can train using a dataset consisting of 1 example per class. DIMCO has this capability because, unlike other methods, DIMCO requires no support/query (also called train/test) split and maximizes the mutual information within a given batch. In contrast, other methods require at least one support and one query example per class within each batch.

For this experiment, we used the Adam optimizer and performed a log-uniform hyperparameter sweep for learning rate $\in [1 \times 10^{-7}, 1 \times 10^{-3}]$. For DIMCO, we swept $p \in [32, 128]$ and $d \in [16, 32]$. For other methods, we made the embedding dimension $\in [16, 32]$. For each combination of loss and number of training examples per class, we ran the experiment 64 times and reported the mean and standard deviation of the top 5.

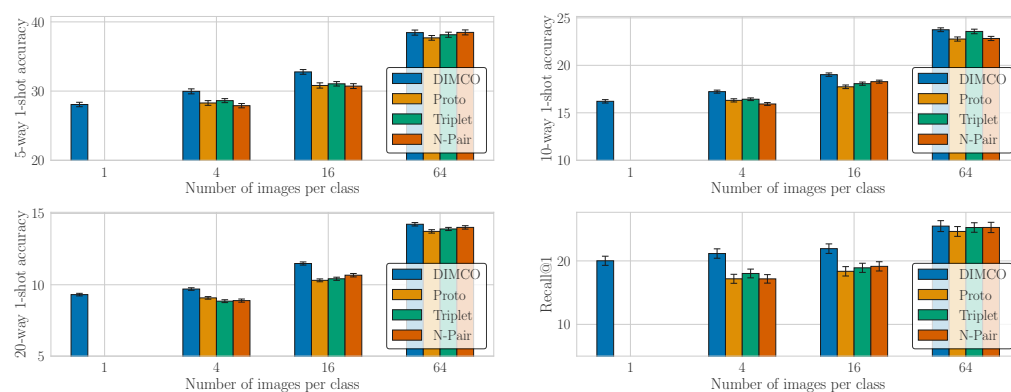


Figure A3. Performance of methods trained using subsets of miniImageNet of varying size. The lowermost y axis value for each metric corresponds to the expected performance of random guessing. We show the mean and standard deviation of top 5 runs from a hyperparameter sweep of 64 runs per configuration.

Appendix C.6. Few-Shot Classification

For this experiment, we built on the code released by Lee et al. [53] (<https://github.com/kjunelee/MetaOptNet>, accessed on 15 March 2021) with minimal adjustments. We used the repository’s default datasets, augmentation, optimizer, and backbones. The only

difference was our added module for outputting discrete codes. We show an extended table with citations in Table A3.

Table A3. Few-shot classification accuracies on the miniImageNet benchmark, with best results for each setting in bold. Grouped according to backbone architecture. † denotes transductive methods, which are more expressive by taking unlabeled examples into account.

Method	5-Way 1-Shot	5-Way 5-Shot
ConvNet (64-64-64-64)		
† TPN [48]	55.51 ± 0.86	69.86 ± 0.65
† FEAT [49]	55.75 ± 0.20	72.17 ± 0.16
MetaLSTM [50]	43.44 ± 0.77	60.60 ± 0.71
MatchingNet [31]	43.56 ± 0.84	55.31 ± 0.73
ProtoNet [3]	49.42 ± 0.78	68.20 ± 0.66
RelationNet [51]	50.44 ± 0.82	65.32 ± 0.70
R2D2 [52]	51.2 ± 0.6	68.8 ± 0.1
MetaOptNet-SVM [53]	52.87 ± 0.57	68.76 ± 0.48
DIMCO (64, 64)	47.33 ± 0.46	61.59 ± 0.52
DIMCO (64, 128)	53.29 ± 0.47	64.79 ± 0.57
ResNet-12		
† TPN [48]	59.46	75.65
† FEAT [49]	62.60 ± 0.20	78.06 ± 0.15
SNAIL [47]	55.71 ± 0.99	68.88 ± 0.92
AdaResNet [57]	56.88 ± 0.62	71.94 ± 0.57
TADAM [4]	58.50 ± 0.30	76.70 ± 0.30
MetaOptNet-SVM [53]	62.64 ± 0.61	78.63 ± 0.46
DIMCO (128, 64)	54.57 ± 0.47	65.45 ± 0.31
DIMCO (128, 128)	57.24 ± 0.44	69.31 ± 0.38

References

- Hoffer, E.; Ailon, N. Deep metric learning using triplet network. In *International Workshop on Similarity-Based Pattern Recognition*; Springer: Berlin/Heidelberg, Germany, 2015; pp. 84–92.
- Movshovitz-Attias, Y.; Toshev, A.; Leung, T.K.; Ioffe, S.; Singh, S. No fuss distance metric learning using proxies. In *Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017*; pp. 360–368.
- Snell, J.; Swersky, K.; Zemel, R. Prototypical networks for few-shot learning. In *Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017*; pp. 4077–4087.
- Oreshkin, B.; López, P.R.; Lacoste, A. Tadam: Task dependent adaptive metric for improved few-shot learning. In *Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018*; pp. 721–731.
- Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. *arXiv* **2000**, arXiv:physics/0004057.
- Shwartz-Ziv, R.; Tishby, N. Opening the black box of deep neural networks via information. *arXiv* **2017**, arXiv:1703.00810.
- Belghazi, M.I.; Baratin, A.; Rajeswar, S.; Ozair, S.; Bengio, Y.; Courville, A.; Hjelm, R.D. MINE: Mutual information neural estimation. In *Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018*.
- Hamming, R.W. Error detecting and error correcting codes. *Bell Syst. Tech. J.* **1950**, *29*, 147–160. [[CrossRef](#)]
- Kuncheva, L.I.; Whitaker, C.J. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Mach. Learn.* **2003**, *51*, 181–207. [[CrossRef](#)]
- Koch, G.; Zemel, R.; Salakhutdinov, R. Siamese neural networks for one-shot image recognition. In *Proceedings of the ICML Deep Learning Workshop, Lille, France, 6–11 July 2015*; Volume 2.
- Sohn, K. Improved deep metric learning with multi-class n-pair loss objective. In *Proceedings of the Advances in Neural Information Processing Systems, Barcelona, Spain, 5–10 December 2016*; pp. 1857–1865.
- Duan, Y.; Zheng, W.; Lin, X.; Lu, J.; Zhou, J. Deep Adversarial Metric Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018*; pp. 2780–2789.
- Amit, R.; Meir, R. Meta-learning by Adjusting Priors Based on Extended PAC-Bayes Theory. In *Proceedings of the International Conference on Machine Learning, Sydney, Australia, 6–11 August 2017*.
- Ravi, S.; Beato, A. Amortized Bayesian Meta-Learning. In *Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019*.
- Shamir, O.; Sabato, S.; Tishby, N. Learning and generalization with the information bottleneck. *Theor. Comput. Sci.* **2010**, *411*, 2696–2711. [[CrossRef](#)]

16. Tishby, N.; Zaslavsky, N. Deep learning and the information bottleneck principle. In Proceedings of the 2015 IEEE Information Theory Workshop (ITW), Jerusalem, Israel, 26 April–1 May 2015; pp. 1–5.
17. Strouse, D.; Schwab, D.J. The deterministic information bottleneck. *Neural Comput.* **2017**, *29*, 1611–1630. [[CrossRef](#)] [[PubMed](#)]
18. Lee, Y.; Choi, S. Gradient-Based Meta-Learning with Learned Layerwise Metric and Subspace. In Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018.
19. Rusu, A.A.; Rao, D.; Sygnowski, J.; Vinyals, O.; Pascanu, R.; Osindero, S.; Hadsell, R. Meta-learning with latent embedding optimization. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
20. Kim, T.; Yoon, J.; Dia, O.; Kim, S.; Bengio, Y.; Ahn, S. Bayesian model-agnostic meta-learning. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018.
21. Zintgraf, L.M.; Shiarlis, K.; Kurin, V.; Hofmann, K.; Whiteson, S. CAML: Fast Context Adaptation via Meta-Learning. In Proceedings of the International Conference on Machine Learning, Long Beach, CA, USA, 10–15 June 2019.
22. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [[CrossRef](#)]
23. Rolfe, J.T. Discrete variational autoencoders. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
24. van den Oord, A.; Vinyals, O. Neural discrete representation learning. In Proceedings of the Advances in Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 6306–6315.
25. Razavi, A.; Oord, A.v.d.; Vinyals, O. Generating Diverse High-Fidelity Images with VQ-VAE-2. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019.
26. Hu, W.; Miyato, T.; Tokui, S.; Matsumoto, E.; Sugiyama, M. Learning discrete representations via information maximizing self-augmented training. In Proceedings of the 34th International Conference on Machine Learning, Sydney, Australia, 6–11 August 2017; Volume 70, pp. 1558–1567.
27. Gray, R.M.; Neuhoff, D.L. Quantization. *IEEE Trans. Inf. Theory* **1998**, *44*, 2325–2383. [[CrossRef](#)]
28. Jegou, H.; Douze, M.; Schmid, C. Product quantization for nearest neighbor search. *IEEE Trans. Pattern Anal. Mach. Intell.* **2011**, *33*, 117–128. [[CrossRef](#)] [[PubMed](#)]
29. Gong, Y.; Lazebnik, S.; Gordo, A.; Perronnin, F. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* **2012**, *35*, 2916–2929. [[CrossRef](#)] [[PubMed](#)]
30. Jeong, Y.; Song, H.O. Efficient end-to-end learning for quantizable representations. In Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018.
31. Vinyals, O.; Blundell, C.; Lillicrap, T.; Wierstra, D. Matching networks for one shot learning. In Proceedings of the Advances in Neural Information Processing Systems, Barcelona, Spain, 5–10 December 2016; pp. 3630–3638.
32. Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In Proceedings of the Advances in Neural Information Processing Systems, Barcelona, Spain, 5–10 December 2016; pp. 2172–2180.
33. Alemi, A.A.; Poole, B.; Fischer, I.; Dillon, J.V.; Saurous, R.A.; Murphy, K. Fixing a Broken ELBO. In Proceedings of the International Conference on Machine Learning, Sydney, Australia, 6–11 August 2017.
34. Chen, T.Q.; Li, X.; Grosse, R.B.; Duvenaud, D.K. Isolating sources of disentanglement in variational autoencoders. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018; pp. 2610–2620.
35. Alemi, A.A.; Fischer, I.; Dillon, J.V.; Murphy, K. Deep Variational Information Bottleneck. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
36. Hjelm, R.D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Trischler, A.; Bengio, Y. Learning deep representations by mutual information estimation and maximization. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
37. Oord, A.v.d.; Li, Y.; Vinyals, O. Representation learning with contrastive predictive coding. *arXiv* **2018**, arXiv:1807.03748.
38. Grover, A.; Ermon, S. Uncertainty autoencoders: Learning compressed representations via variational information maximization. In Proceedings of the International Conference on Machine Learning, Long Beach, CA, USA, 10–15 June 2019.
39. Choi, K.; Tatwawadi, K.; Grover, A.; Weissman, T.; Ermon, S. Neural Joint Source-Channel Coding. In Proceedings of the International Conference on Machine Learning, Long Beach, CA, USA, 10–15 June 2019; pp. 1182–1192.
40. Achille, A.; Soatto, S. Information dropout: Learning optimal representations through noisy computation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *40*, 2897–2905. [[CrossRef](#)] [[PubMed](#)]
41. Kim, H.; Mnih, A. Disentangling by factorising. In Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018; pp. 2649–2658.
42. Krizhevsky, A. *Learning Multiple Layers of Features from Tiny Images*; Technical Report; University of Tront: Toronto, ON, Canada, 2009.
43. Wah, C.; Branson, S.; Welinder, P.; Perona, P.; Belongie, S. *The Caltech-UCSD Birds-200-2011 Dataset*; Technical Report; California Institute of Technology: Pasadena, CA, USA, 2011.
44. Krause, J.; Stark, M.; Deng, J.; Fei-Fei, L. 3D Object Representations for Fine-Grained Categorization. In Proceedings of the 4th International IEEE Workshop on 3D Representation and Recognition (3dRR-13), Sydney, Australia, 2–8 December 2013.

45. Deng, J.; Dong, W.; Socher, R.; Li, L.J.; Li, K.; Fei-Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL, USA, 20–25 June 2009.
46. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HA, USA, 21–26 July 2016; pp. 770–778.
47. Mishra, N.; Rohaninejad, M.; Chen, X.; Abbeel, P. A simple neural attentive meta-learner. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
48. Liu, Y.; Lee, J.; Park, M.; Kim, S.; Yang, Y. Transductive propagation network for few-shot learning. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
49. Ye, H.J.; Hu, H.; Zhan, D.C.; Sha, F. Learning embedding adaptation for few-shot learning. *arXiv* **2018**, arXiv:1812.03664.
50. Ravi, S.; Larochelle, H. Optimization as a model for few-shot learning. In Proceedings of the International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
51. Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P.H.; Hospedales, T.M. Learning to compare: Relation network for few-shot learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 1199–1208.
52. Bertinetto, L.; Henriques, J.F.; Torr, P.H.; Vedaldi, A. Meta-learning with differentiable closed-form solvers. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
53. Lee, K.; Maji, S.; Ravichandran, A.; Soatto, S. Meta-learning with differentiable convex optimization. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15–20 June 2019; pp. 10657–10665.
54. Yi, D.; Lei, Z.; Liao, S.; Li, S.Z. Deep metric learning for person re-identification. In Proceedings of the 2014 22nd International Conference on Pattern Recognition, Stockholm, Sweden, 24–28 August 2014; pp. 34–39.
55. Yin, M.; Tucker, G.; Zhou, M.; Levine, S.; Finn, C. Meta-Learning without Memorization. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
56. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. In Proceedings of the International Conference on Learning Representations, San Diego, CA, USA, 7–9 May 2015.
57. Munkhdalai, T.; Yuan, X.; Mehri, S.; Trischler, A. Rapid adaptation with conditionally shifted neurons. In Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018.