

EFFICIENT DISTRIBUTION MATCHING OF REPRESENTATIONS VIA NOISE-INJECTED DEEP INFOMAX

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep InfoMax (DIM) is a well-established method for self-supervised representation learning (SSRL) based on maximization of the mutual information between the input and the output of a deep neural network encoder. Despite the DIM and contrastive SSRL in general being well-explored, the task of learning representations conforming to a specific distribution (i.e., distribution matching, DM) is still under-addressed. Motivated by the importance of DM to several downstream tasks (including generative modeling, disentanglement, outliers detection and other), we enhance DIM to enable automatic matching of learned representations to a selected prior distribution. To achieve this, we propose injecting an independent noise into the normalized outputs of the encoder, while keeping the same InfoMax training objective. We show that such modification allows for learning uniformly and normally distributed representations, as well as representations of other absolutely continuous distributions. Our approach is tested on various downstream tasks. The results indicate a moderate trade-off between the performance on the downstream tasks and quality of DM.

1 INTRODUCTION

Learning viable low-dimensional representations of complex data plays an important role in many modern applications of artificial intelligence. This task arises in various domains, including image (Haralick et al., 1973; Chen et al., 2020b; Rombach et al., 2022), audio (van den Oord et al., 2019), and natural language processing (Mikolov et al., 2013; Radford et al., 2018; Devlin et al., 2019). High-quality embeddings are particularly useful for multi-modal methods (Vinyals et al., 2015; Radford et al., 2021; Ho & Salimans, 2021), statistical and topological analysis (Moor et al., 2020; Duong & Nguyen, 2022; Butakov et al., 2024b), data visualization (van der Maaten & Hinton, 2008; McInnes et al., 2018), and testing fundamental hypotheses (Brown et al., 2023; Gurnee & Tegmark, 2024; Huh et al., 2024).

Existing approaches to representation learning can be divided into three categories (Ericsson et al., 2022): *supervised* (requires labeled data), *self-supervised* and *unsupervised* (no labeling is required). In practice, access to labeled data is limited, which hinders the use of supervised approaches. Therefore, unsupervised and self-supervised methods are of great importance. Contrastive learning is a well-established paradigm of self-supervised representation learning (SSRL), which encourages an encoder to learn similar representations for various augmentations of the same data point, and dissimilar – for different data points. Deep InfoMax (DIM) (Hjelm et al., 2019) leverages information-theoretic quantities to construct a decent contrastive objective, involving a direct maximization of the useful information contained in the embeddings. DIM is universal, flexible, and deeply connected to the rigorous information theory, which allows for a good performance on a variety of downstream tasks to be attained (Hjelm et al., 2019; Bachman et al., 2019; Veličković et al., 2019; Tschannen et al., 2020; Yu, 2024).

Acquiring embeddings admitting a specific distribution (i.e., distribution matching, DM) is an auxiliary, yet important task in representation learning. Latent distributions with straightforward sampling procedures or tractable densities are crucial for downstream generative modeling (Kingma & Welling, 2014; Makhzani et al., 2016; Larsen et al., 2016; Papamakarios et al., 2021). Additionally, specific distributions (e.g., Gaussian) exhibit properties, which are useful for statistical analysis (Tip-

ping & Bishop, 1999; Duong & Nguyen, 2022), disentanglement (Higgins et al., 2017; Balabin et al., 2024), and outliers detection.

A classical approach to latent DM is to optimize a cheap and imprecise distribution dissimilarity measure during the training or architecture search (Ng, 2011; Makhzani & Frey, 2014; Kingma & Welling, 2014; Heusel et al., 2017; Higgins et al., 2017). Methods of this family have to rely on several strong assumptions, such as embeddings already admitting a Gaussian distribution, which eventually leads to suboptimal results. Another common approach employs adversarial networks to push the learned representations towards a desired distribution (Makhzani et al., 2016; Hjelm et al., 2019). One can also leverage generative models (e.g., normalizing flows or diffusion models) to explicitly perform DM “post hoc” (Böhm & Seljak, 2022; Rombach et al., 2022). These two approaches yield decent results, but require supplementary networks to match the distributions. Finally, injective likelihood-based models can also be used to acquire low-dimensional representations admitting a desired distribution (Brehmer & Cranmer, 2020; Sorrenson et al., 2024). However, likelihood maximization across dimensions is notoriously problematic due to non-square Jacobi matrices.

In contrast to the mentioned approaches, we propose a simple, cost-effective and non-intrusive modification to DIM, which allows for an automatic and exact DM of the representations. In the following text, we show that using specific activation functions and noise injections at the outlet of an encoder, combined with the DIM objective, allows for normally and uniformly distributed representations to be learned. Our contributions are the following:

1. We prove that applying normalization and adding a small noise at the outlet of an encoder and maximizing the DIM objective minimizes the Kullback-Leibler divergence between a Gaussian (or uniform) distribution and the distribution of embeddings.
2. We conduct experiments on several downstream tasks to explore the trade-off between the downstream performance and the accuracy of DM via our method.
3. We conduct additional experiments to assess the quality of DM via our method in the tasks of generative modelling.

The paper is organized as follows. In Section 2, the necessary background from information theory and original works on Deep InfoMax is provided. In Section 3, we describe the general method for DM via modified DIM. In Section 4, a connection between the proposed approach and other SSRL methods is established. Section 5 is dedicated to the experimental evaluation of our method. Finally, we conclude the paper by discussing our results in Section 6. Complete proofs, additional analysis of infomax-based DM, and technical details are provided in Appendices A to C correspondingly.

2 BACKGROUND

In this section, the background necessary to understand our work is provided. We start with the basic definitions from the information theory. Then, the maximum entropy theorems are given, which are crucial for understanding our approach. Finally, we outline the general variant of Deep InfoMax representation learning method, which we aim to enhance with an automatic distribution matching.

2.1 PRELIMINARIES

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space with sample space Ω , σ -algebra \mathcal{F} , and probability measure \mathbb{P} defined on \mathcal{F} . Consider an absolutely continuous random vector $X: \Omega \rightarrow \mathbb{R}^d$ with the probability density function (PDF) denoted as $p(x)$. The differential entropy of X is defined as follows:

$$h(X) = -\mathbb{E} \log p(X) = -\int_{\text{supp } X} p(x) \log p(x) dx,$$

where $\text{supp } X \subseteq \mathbb{R}^n$ represents the *support* of X , and $\log(\cdot)$ denotes the natural logarithm. Similarly, we define the joint differential entropy as $h(X, Y) = -\mathbb{E} \log p(x, y)$ and conditional differential entropy as $h(X | Y) = -\mathbb{E} \log p(X|Y) = -\mathbb{E}_Y (\mathbb{E}_{X|Y} \log p(X | Y))$. Finally, the mutual information (MI) is given by $I(X; Y) = h(X) - h(X | Y)$, and the following equivalences hold

$$I(X; Y) = h(X) - h(X | Y) = h(Y) - h(Y | X),$$

$$I(X; Y) = h(X) + h(Y) - h(X, Y),$$

$$I(X; Y) = D_{\text{KL}}(\mathbb{P}_{X,Y} \parallel \mathbb{P}_X \otimes \mathbb{P}_Y).$$

Mutual information can also be defined as an expectation of the *pointwise mutual information*:

$$\text{PMI}_{X,Y}(x, y) = \log \left[\frac{p(x | y)}{p(x)} \right], \quad I(X; Y) = \mathbb{E} \text{PMI}_{X,Y}(X, Y). \quad (1)$$

The above definitions can be generalized via Radon-Nikodym derivatives and induced densities in case of distributions supports being manifolds, see (Spivak, 1965). We use flexible notation, where entropy and divergence may refer to random vectors or their corresponding distributions.

Our work leverages the maximum entropy properties of Gaussian and uniform distributions:

Theorem 2.1 (Theorem 8.6.5 in Cover & Thomas (2006)). *Let X be a d -dimensional absolutely continuous random vector with probability density function p , mean m and covariance matrix Σ . Then*

$$h(X) = h(\mathcal{N}(m, \Sigma)) - D_{\text{KL}}(p \parallel \mathcal{N}(m, \Sigma)), \quad h(\mathcal{N}(m, \Sigma)) = \frac{1}{2} \log((2\pi e)^d \det \Sigma),$$

where $\mathcal{N}(m, \Sigma)$ is a Gaussian distribution of mean m and covariance matrix Σ .

Theorem 2.2. *Let X be an absolutely continuous random vector with probability density function p and $\text{supp } X \subseteq S$, where S has finite and non-zero Lebesgue measure $\mu(S)$. Then*

$$h(X) = h(\text{U}(S)) - D_{\text{KL}}(p \parallel \text{U}(S)), \quad h(\text{U}(S)) = \log \mu(S),$$

where $\text{U}(S)$ is a uniform distribution on S .

Remark 2.3. Note that $h(X) \neq h(X') - D_{\text{KL}}(p_X \parallel p_{X'})$ in general.

Finally, we also utilize the following lower bound on the conditional entropy of a sum of two conditionally independent random vectors:

Lemma 2.4. *Let X and Z be random vectors of the same dimensionality, independent under the conditioning vector Y . Then*

$$h(X + Z | Y) = h(Z | Y) + I(X; X + Z | Y) \geq h(Z | Y),$$

with equality if and only if there exists a measurable function g such that $X = g(Y)$.¹

2.2 DEEP INFOMAX

Mutual information (MI) is widely considered as a fundamental measure of statistical dependence between random variables due to its key properties, such as invariance under diffeomorphisms, sub-additivity, non-negativity, and symmetry (Cover & Thomas, 2006). These attributes make MI particularly useful in information-theoretic approaches to machine learning. The concept of maximizing MI between input and output, known as the infomax principle (Linsker, 1988; Bell & Sejnowski, 1995), serves as the foundation for Deep InfoMax (DIM), a family of self-supervised representation learning methods proposed by Hjelm et al. (2019).

A naïve DIM approach suggests learning the most informative embeddings via a direct maximization of the mutual information between the original data and compressed representations:

$$I(X; f(X)) \rightarrow \max,$$

where X is the random vector to be compressed, and f is the encoding mapping (being learned). However, despite its simplicity and intuitiveness, this setting is rendered useless by the fact that $I(X; f(X)) = \infty$ for a wide range of X and f (Bell & Sejnowski, 1995; Hjelm et al., 2019).

To avoid this limitation, it is usually suggested to augment (crop randomly, add noise, etc.) original data to inject stochasticity and make information-theoretic objectives non-degenerate. Consider the following Markov chain:

$$f(X) \longrightarrow X \longrightarrow X',$$

¹Hereinafter, when comparing random variables, we mean equality or inequality “almost sure”.

where X' is augmented data. Now $I(X'; f(X))$ can be made non-infinite. Moreover, according to the data processing inequality (Cover & Thomas, 2006), $I(X'; f(X)) \leq I(X; f(X))$, which connects this new infomax objective with the naïve one.

The only remaining problem is the high dimensionality of X' , which makes MI estimation difficult. To resolve this, one can apply an additional (in general, random) dimensionality-reducing transformation to replace X' with Y' :

$$f(X) \longrightarrow X \longrightarrow X' \longrightarrow Y',$$

with $I(Y'; f(X)) \leq I(X'; f(X))$ being the new infomax objective. For the sake of computational simplicity, Y' can be defined as $f(X')$; thus, the same encoder network is used to compress both the original and augmented data. Moreover, one can view $I(f(X'); f(X))$ as a contrastive objective: this value is high when f yields similar representations for corresponding augmented and non-augmented samples, and dissimilar for other pairs of samples.

Now, let us consider the task of mutual information maximization. As MI is notoriously hard to estimate (McAllester & Stratos, 2020), lower bounds are widely used to reparametrize the infomax objective (Belghazi et al., 2021; van den Oord et al., 2019; Hjelm et al., 2019). In this work, we only consider the two most popular approaches based on the variational representations of the Kullback-Leibler divergence:

$$\begin{array}{l} \textit{Donsker-Varadhan} \\ \text{(Donsker \& Varadhan 1983)} \\ \text{Belghazi et al. 2021)} \end{array} \quad I(X; Y) = \sup_{T: \Omega \rightarrow \mathbb{R}} [\mathbb{E}_{\mathbb{P}_{X,Y}} T - \log \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Y} \exp(T)] \quad (2)$$

$$\begin{array}{l} \textit{Nguyen-Wainwright-Jordan} \\ \text{(Nguyen et al. 2010,)} \\ \text{Belghazi et al. 2021)} \end{array} \quad I(X; Y) = \sup_{T: \Omega \rightarrow \mathbb{R}} [\mathbb{E}_{\mathbb{P}_{X,Y}} T - \mathbb{E}_{\mathbb{P}_X \otimes \mathbb{P}_Y} \exp(T - 1)] \quad (3)$$

where Ω is the sampling space, and T is a measurable critic function.

Remark 2.5. Assuming $\text{PMI}_{X,Y}$ exists, the supremum in (2) and (3) is attained at and only at

$$T^* = T^*(x, y) = \text{PMI}_{X,Y}(x, y) + \alpha \quad (\text{a.s. w.r.t. } \mathbb{P}_{X,Y}), \quad (4)$$

where $\alpha = 1$ for (3) and can be any real number for (2) (follows from Theorem 1 in Belghazi et al. (2021) and Theorem 2.1 in Keziou (2003)).

In practice, f and T are approximated via corresponding neural networks, with the parameters being learned through the maximization of the Monte-Carlo estimate of (2) or (3). Although (2) usually yields better estimates, the SGD gradients of this expression are biased in a mini-batch setting (Belghazi et al., 2021).

3 DEEP INFOMAX WITH AUTOMATIC DISTRIBUTION MATCHING

In the present section, we modify DIM to enable automatic distribution matching (DM) of representations. Specifically, we propose adding independent noise Z to normalized representations of X , thus replacing $f(X)$ by $f(X) + Z$ in the infomax objective from Section 2.2. This corresponds to the following Markov chain:

$$f(X) + Z \longrightarrow f(X) \longrightarrow X \longrightarrow X' \longrightarrow f(X').$$

By the data processing inequality, we know that $I(f(X'); f(X) + Z) \leq I(f(X'); f(X))$, which connects our method to the family of conventional DIM approaches discussed previously.

As we will demonstrate, the noise injection at the outlet of an encoder, in conjunction with Theorems 2.1 and 2.2 on maximum entropy, allows us to achieve normally or uniformly distributed embeddings. To prove this, we use the following Lemma 3.1, which decomposes the proposed infomax objective into three components: the entropy of representations with an additive noise Z , the entropy of the noise itself, and the mutual information between $f(X)$ and $f(X) + Z$ conditioned on $f(X')$ – representations of augmented data.

Lemma 3.1. *Consider the following Markov chain of absolutely continuous random vectors:*

$$f(X) + Z \longrightarrow X \longrightarrow X' \longrightarrow f(X'),$$

with Z being independent of (X, X') . Then

$$I(f(X'); f(X) + Z) = h(f(X) + Z) - h(Z) - I(f(X) + Z; f(X) | f(X')). \quad (5)$$

Lemma 3.1 highlights the importance of both additive noise and input data augmentation. Specifically, if no augmentation is applied, i.e., $X = X'$, then $I(f(X) + Z; f(X) | f(X')) = 0$. In this case, maximizing the infomax objective is reduced to maximizing the entropy $h(f(X) + Z)$. Under the corresponding restrictions on $f(X)$, this is equivalent to distribution matching, see Theorems 2.1 and 2.2. In contrast, if no noise is added, $I(f(X'); f(X))$ is not bounded in general, which may prevent $h(f(X))$ from saturation in practical scenarios.

Additionally, DM alone does not guarantee that the learned representations will be meaningful or useful for downstream tasks. That is why we also show that using $X \neq X'$ allows us to recover meaningful embeddings, as $I(f(X) + Z; f(X) | f(X'))$ is set to zero by learning representations that are *weakly invariant* to selected data augmentations:

Definition 3.2. We call an encoding mapping f *weakly invariant* to data augmentation $X \rightarrow X'$ if there exists a function g such that $f(X) = g(f(X)) = g(f(X'))$ almost surely.

Lemma 3.3. *Under the conditions of Lemma 3.1, let $\mathbb{P}(X = X' | X) \geq \alpha > 0$. Then, $I(f(X) + Z; f(X) | f(X')) = 0$ precisely when f is weakly invariant to $X \rightarrow X'$.*

Overall, the signal-to-noise ratio serves as a tradeoff between the distribution matching objective and robustness to the data augmentations (the first and the third term in (5) correspondingly): higher magnitudes of Z impose tighter bounds on $I(f(X) + Z; f(X) | f(X'))$, thus prioritizing maximization of $h(f(X) + Z)$. In what follows, we formalize the provided reasoning by addressing Gaussian and uniform distribution matching separately. We also briefly discuss how a general DM problem can be reduced to the normal or uniform one.

3.1 GAUSSIAN DISTRIBUTION MATCHING

Let us assume Z having finite second-order moments. According to Theorem 2.1, if we restrict $f(X)$ by fixing its covariance matrix, $h(f(X) + Z)$ attains its maximal value precisely when $f(X) + Z$ is distributed normally. Now, as Z is independent of X , $f(X) + Z$ is normally distributed if and only if the distribution of $f(X)$ is also Gaussian.

Thus, one can achieve the normality of $f(X)$ via restricting second-order moments of $f(X)$ and maximizing the entropy $h(f(X) + Z)$. In combination with Lemma 3.1, this forms a basis of the proposed method for Gaussian distribution matching. Moreover, we show that the imposed restrictions can be partially lifted via only requiring $\text{Var}(f(X)_i) = 1$ for every $i \in \{1, \dots, d\}$. This approach is preferable in practice, as the widely-used batch normalization (Ioffe & Szegedy, 2015) can be employed to restrict the variances. We formalize our findings in the following theorem.

Theorem 3.4 (Gaussian distribution matching). *Let the conditions of Lemma 3.3 be satisfied. Assume $Z \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, $\mathbb{E} f(X) = 0$ and $\text{Var}(f(X)_i) = 1$ for all $i \in \{1, \dots, d\}$. Then, the mutual information $I(f(X'); f(X) + Z)$ can be upper bounded as follows*

$$I(f(X'); f(X) + Z) \leq \frac{d}{2} \log \left(1 + \frac{1}{\sigma^2} \right), \quad (6)$$

with the equality holding exactly when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$. Moreover,

$$D_{\text{KL}}(f(X) \| \mathcal{N}(0, \mathbf{I})) \leq I(Z; f(X) + Z) - I(f(X'); f(X) + Z) - d \log \sigma.$$

While the first inequality in Theorem 3.4 suggests that exact distribution matching is achieved only asymptotically, the second expression shows that $D_{\text{KL}}(f(X) \| \mathcal{N}(0, \mathbf{I}))$ can be bounded using the value of the proposed infomax objective. Note that $I(Z; f(X) + Z)$ is typically upper bounded, except the rare cases where the support of $f(X)$ is degenerate.

3.2 UNIFORM DISTRIBUTION MATCHING

Now consider Z distributed uniformly on $[-\varepsilon; \varepsilon]^d$, and let $f(X)$ has bounded support on $[0; 1]^d$. By Theorem 2.2, the entropy $h(f(X) + Z)$ reaches its maximum precisely when $f(X) + Z$ is uniformly distributed. While the independence of Z and $f(X)$ does not imply that $f(X)$ is uniform, the distribution of $f(X)$ that maximizes $h(f(X) + Z)$ is a discrete uniform distribution over a finite grid within $[0; 1]^d$. As ε approaches zero, the distribution of $f(X)$ converges to a continuous uniform distribution over $[0; 1]^d$. Thus, injecting uniform noise asymptotically drives $f(X)$ towards a uniform representation over $[0; 1]^d$. This reasoning is formalized in the following theorem.

Theorem 3.5 (Uniform distribution matching). *Under the conditions of Lemma 3.1, let $Z \sim \mathcal{U}([- \varepsilon; \varepsilon]^d)$ and $\text{supp } f(X) \subseteq [0; 1]^d$. Then, the mutual information $I(f(X'); f(X) + Z)$ can be upper bounded as follows*

$$I(f(X'); f(X) + Z) \leq d \log \left(1 + \frac{1}{2\varepsilon} \right). \quad (7)$$

with the equality if and only if $1/\varepsilon \in \mathbb{N}$, f is weakly invariant, and $f(X) \sim \mathcal{U}(A)$, where the set $A = \{0, 2\varepsilon, 4\varepsilon, \dots, 1\}$ contains $(1/(2\varepsilon) + 1)$ elements.

Moreover,

$$D_{\text{KL}}(f(X) \parallel \mathcal{U}([0; 1]^d)) \leq I(Z; f(X) + Z) - I(f(X'); f(X) + Z) - d \log(2\varepsilon).$$

In sharp contrast to Theorem 3.4, the equality in (7) is attained at $f(X)$ conforming to a discrete distribution. This makes the proposed objective less attractive in comparison to the Gaussian distribution matching. However, one still can be assured that the standard continuous uniform distribution allows us to approach the equality in (7) with ε approaching zero:

Remark 3.6 (Butakov et al. 2024a). If $f(X) \sim \mathcal{U}([0; 1]^d)$, f is weakly invariant, and $\varepsilon < 1/2$, then

$$I(f(X'); f(X) + Z) = d(\varepsilon - \log(2\varepsilon)) = d \log(1 + 1/(2\varepsilon)) - \underbrace{(\log(1 + 2\varepsilon) - \varepsilon)}_{o(\varepsilon)}.$$

3.3 GENERAL CASE

Our approach can be extended to a wide range of desired distributions of embeddings using the probability integral transform (David & Johnson, 1948; Chen & Gopinath, 2000). Normalizing flows (learnable diffeomorphisms) can also be leveraged to transform the distribution “post-hoc” due to their universality property (Huang et al., 2018; Jaini et al., 2019; Kobyzev et al., 2020).

The data processing inequality can be used to estimate the quality of DM after the transformation:

Statement 3.7 (Corollary 2.18 in (Polyanskiy & Wu, 2024)). *Let g be a measurable (w.r.t \mathbb{P} and \mathbb{Q}) function. Then $D_{\text{KL}}(\mathbb{P} \circ g^{-1} \parallel \mathbb{Q} \circ g^{-1}) \leq D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q})$, where $\mathbb{P} \circ g^{-1}$ and $\mathbb{Q} \circ g^{-1}$ denote the push-forward measures of \mathbb{P} and \mathbb{Q} after applying g .*

4 CONNECTION TO OTHER METHODS FOR SSRL

In this section, we explore the relation of our approach to other methods for unsupervised and self-supervised representation learning. We argue that the proposed technique for distribution matching can be applied to any other method, given the latter can be formulated in terms of mutual information maximization.

Autoencoders In reconstruction-based generative models (Makhzani et al., 2016), the reconstruction error is closely tied to the mutual information (Theorem 8.6.6 in (Cover & Thomas, 2006)):

$$I(X; Y) = h(X) - h(X | Y) \geq h(X) - \frac{d}{2} \log(2\pi e) - \frac{1}{2} \log \mathbb{E}[\|X - \hat{X}(Y)\|^2].$$

Here, X denotes the input, Y the latent representation produced by the encoder, and $\hat{X}(Y)$ the reconstructed input. The last term involving the expected squared error² is essentially the autoencoder’s loss function. By minimizing this reconstruction loss, the mutual information between the input data and its representation is maximized. However, recall that $I(X; Y)$ can diverge to infinity, as discussed earlier, so it is crucial to introduce augmentations.

InfoNCE Many conventional methods for self-supervised representation learning leverage InfoNCE loss (van den Oord et al., 2019) in conjunction with a similarity measure $T(\cdot, \cdot)$ to formulate the following contrastive objective (Tschannen et al., 2020; He et al., 2020; Chen et al., 2020a):

$$\mathcal{L}_{\text{InfoNCE}} = -[\mathbb{E}_{\mathbb{P}^+} T(Q, K) - \log \mathbb{E}_{\mathbb{P}^-} \exp(T(Q, K))], \quad \hat{\mathcal{L}}_{\text{InfoNCE}} = -\log \frac{e^{T(q, k^+)}}{\frac{1}{K} \sum_{i=1}^K e^{T(q, k_i^-)}},$$

²Here and throughout, $\|\cdot\|$ denotes the Euclidean norm.

where \mathbb{P}^+ and \mathbb{P}^- denote the distributions of positive and negative pairs of keys K and queries Q correspondingly. By selecting $\mathbb{P}^+ = \mathbb{P}_{f(X'), f(X)+Z}$ and $\mathbb{P}^- = \mathbb{P}_{f(X') \otimes \mathbb{P}_{f(X)+Z}}$, we recover the Donsker-Varadhan bound (2) for our infomax objective $I(f(X'); f(X) + Z)$. However, note that in (2) and (3) the supremum is taken over all measurable functions. In contrast, non-infomax methods typically employ a *separable critic*: $T(q, k) = \langle \phi(q), \psi(k) \rangle$, where ϕ and ψ are *projection heads* (Tschannen et al., 2020; Chen et al., 2020a). In special cases $\phi, \psi = \text{Id}$, so similarity of representations is measured via a plain dot product.

Despite this significant difference, our distribution matching paradigm allows us to drop the supremum in (2) and (3) and establish a direct connection between DIM and traditional non-infomax contrastive SSRL methods:

Theorem 4.1 (Dual form of Gaussian distribution matching). *Under the conditions of Theorem 3.4,*

$$I(f(X'); f(X) + Z) \geq \mathbb{E}_{\mathbb{P}^+} \left[T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^* \right] - \log \mathbb{E}_{\mathbb{P}^-} \left[\exp \left(T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^* \right) \right],$$

$$T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^*(x, y) = \frac{\|y\|^2}{2(1 + \sigma^2)} - \frac{\|y - x\|^2}{2\sigma^2} = \frac{1}{\sigma^2} \left(\langle x, y \rangle - \frac{\|x\|^2 + \|y\|^2 / (1 + \sigma^2)}{2} \right),$$

with the equality holding precisely when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$.

Note that $\langle x, y \rangle$ is widely used as a similarity measure, σ^2 can be interpreted as temperature, and the remaining part of the expression serves as a regularization term.

Covariance-based methods Barlow Twins (Zbontar et al., 2021) and VICReg (Bardes et al., 2021) objective functions can be traced to information-theoretic terms either through Information Bottleneck (Tishby et al., 1999) or multi-view InfoMax (Federici et al., 2020) principles:

$$h(f(X') | X) - \lambda h(f(X')) \rightarrow \min \quad (\text{BarlowTwins})$$

$$I(f(X'); X') \geq h(f(X')) + \mathbb{E}[\log q(f(X') | X'')] \rightarrow \max \quad (\text{VICReg})$$

where $\lambda > 0$, $X \rightarrow X''$ is a separate augmentation path, independent of $X \rightarrow X'$, and q is a probability density function of some distribution. In both cases, the representation entropy $h(f(X'))$ comes into play. Recall that normal distribution is the maximum entropy distribution given first two moments (Theorem 2.1). As both methods employ covariance restriction terms in the respective objective functions, one can recover these methods by assuming the normality of $f(X')$.

5 EXPERIMENTS

In this section, we evaluate our distribution matching approach on several datasets and downstream tasks. To assess the quality of the embeddings, we solve downstream classification tasks and calculate clustering scores. To explore the relation between the magnitude of injected noise and the quality of DM, a set of statistical normality tests is employed. For the experiments requiring numerous evaluations or visualization, we use MNIST handwritten digits dataset LeCun et al. (2010). For other experiments, we use CIFAR10 and CIFAR100 datasets (Krizhevsky, 2009).

Multivariate normality and uniformity tests The key part of the experimental pipeline is to estimate how much the distribution of embeddings acquired via the proposed method is similar to the multivariate normal or uniform distribution. To do this, we leverage D’Agostino-Pearson (D’Agostino, 1971; D’Agostino & Pearson, 1973), Shapiro-Wilk González-Estrada et al. (2022) univariate tests, and Henze-Zirkler (Henze & Zirkler, 1990; Trujillo-Ortiz et al., 2007) multivariate test. To extend the univariate tests to higher dimensions, we utilize the fundamental property of the Gaussian distribution, which is the normality of any linear projection.

In practice, we sample random projectors from a d -dimensional sphere to perform a univariate normality test. We also employ bootstrapping with small subsampling size to get low-variance averaged p -values and to smooth the transition from low to high p -values. Finally, we reduce the uniform distribution case to the Gaussian via the probability integral transform, see Section 3.3. We report the results of the tests in Figure 1. We also visualize the two-dimensional embeddings of the MNIST dataset in Figure 2.

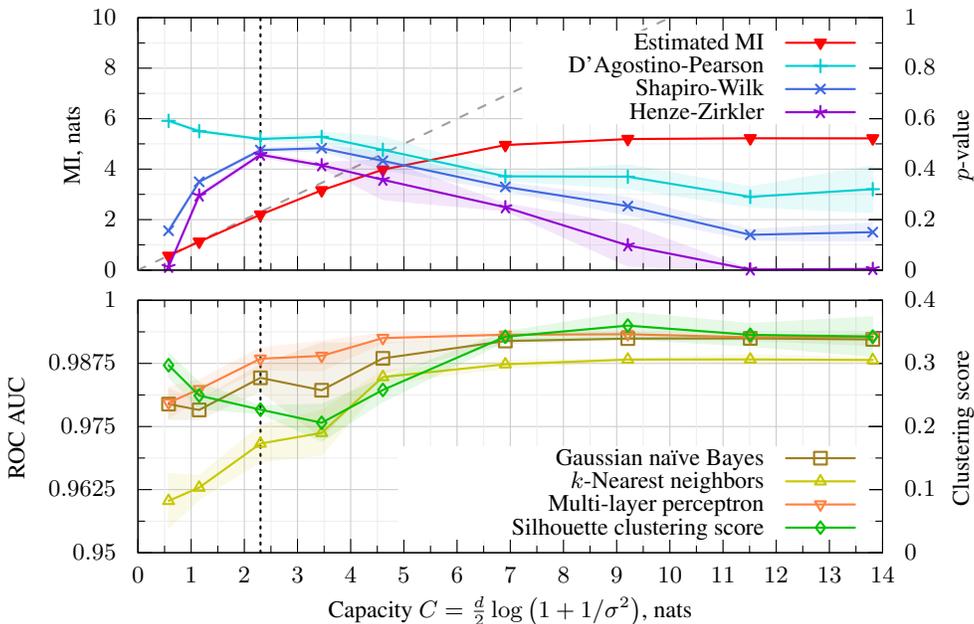


Figure 1: Results for MNIST dataset in the Gaussian DM setup for $d = 2$ with varying capacity $C = \frac{d}{2} \log(1 + 1/\sigma^2)$, measured in nats (units of information based on natural logarithms). The dotted line denotes the minimal capacity required to preserve the information about the class labels in $f(X) + Z$. The dashed line represents the upper bound on the mutual information (6). We run 5 experiments for each point and report mean values and 99% asymptotic confidence intervals. InfoNCE loss is used to approximate (2).

Classification and clustering To explore the trade-off between the magnitude of the injected noise and the quality of representations, we evaluate various clustering metrics and perform downstream classification using conventional ML methods, such as Gaussian naïve Bayes, k -nearest neighbors and shallow multilayer perceptron. In order to numerically measure the quality of clustering, we compute Silhouette score (Rousseeuw, 1987). The results are reported in Figure 1.

We also verify noise injection does not affect typical methods for SSRL such as SimCLR (Chen et al., 2020a) and VICReg (Bardes et al., 2021). To this end, we train both methods on CIFAR-10 and CIFAR-100 datasets with varying degree of noise. The linear probing performance does not drop when noise magnitude is increased across σ values (0.0, 0.1, 0.3, 0.5) as seen in Table 1.

	CIFAR-10		CIFAR-100	
	top-1	top-5	top-1	top-5
SimCLR	90.83	99.76	65.64	89.91
SimCLR $\sigma = 0.1$	90.96	99.72	67.03	90.49
SimCLR $\sigma = 0.3$	91.56	99.77	65.72	89.76
SimCLR $\sigma = 0.5$	90.51	99.74	65.58	89.56
VICReg	90.63	99.67	65.71	88.96
VICReg $\sigma = 0.1$	91.09	99.68	68.92	90.50
VICReg $\sigma = 0.3$	90.75	99.61	67.31	89.89
VICReg $\sigma = 0.5$	91.02	99.75	66.52	89.66

Table 1: Noise magnitude influence on linear probing on CIFAR-10 and CIFAR-100. Top-1 and top-5 accuracy (in %) is reported.

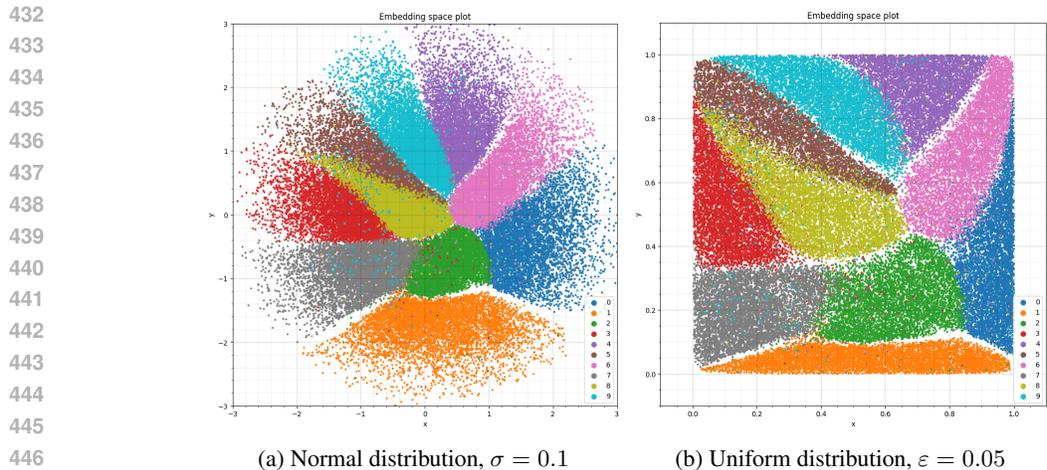


Figure 2: Visualization of two-dimensional representations of the MNIST handwritten digits dataset.

Generation Generative adversarial networks (Goodfellow et al., 2014) produce samples from the underlying latent distribution, which is usually Gaussian. A common approach to make the generation not purely random is to introduce a conditioning vector during the generation (Mirza & Osindero, 2014; Larsen et al., 2016). However, the distribution of the conditioning vector is usually unknown except for specific cases, which hinders the unconditional (or partially conditional) generation via the same model. Thus, a model converting data to the corresponding conditioning vectors with known prior distribution is of great use. In Appendix D, we plug our encoder into a conditional GAN and perform conditional and unconditional generation.

6 DISCUSSION

In this paper, we have proposed and thoroughly investigated a novel and efficient approach to the problem of distribution matching of learned representations. Our technique falls into a well-established family of Deep InfoMax self-supervised representation learning methods, and does not require solving min-max optimization problems or employing generative models “post-hoc” to match the distributions. The proposed approach is grounded in the information theory, which allows for a rigorous theoretical justification of the method. In our work, we also explore the possibility of applying our technique to other popular methods for unsupervised and self-supervised representation learning. Consequently, we assert that the proposed approach can be utilized in conjunction with any other method for SSRL, provided that it can be formulated in terms of mutual information or entropy maximization.

To assess the quality of the representations yielded by our method, we (a) visualize the embeddings and run normality tests, and (b) solve a set of downstream tasks in various experimental setups. The results indicate the following:

1. Increasing the noise magnitude facilitates better distribution matching, but only up to a certain point; beyond that, the entire representation learning process begins to deteriorate due to insufficient information being transmitted through the noisy channel.
2. Experiments with MNIST and low-dimensional embeddings indicate the existence of a moderate trade-off between the magnitude of the injected noise and the quality in downstream classification tasks. However, experiments with other infomax-related methods for SSRL and higher embedding dimensionality suggest this influence being negligible.
3. Embeddings acquired via the proposed method can be used to condition generative models, allowing both conditional and unconditional generation due to the distribution of the conditioning vector being known.

486 **Future work** As for further research, we consider elaborating on the dual formulation of infomax-
 487 based distribution matching. Theorem 4.1 suggests that Gaussian DM can be achieved through a
 488 specific choice of the critic network $T(x, y)$. We plan extending this result to other distributions.
 489 Additionally, we consider conducting further experiments with generative models to comprehen-
 490 sively assess the quality of conditional and unconditional generation, utilizing conditioning vectors
 491 obtained via the proposed method.

492 REFERENCES

- 493 Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing
 494 mutual information across views. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc,
 495 E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32.
 496 Curran Associates, Inc., 2019. URL [https://proceedings.neurips.cc/paper_
 497 files/paper/2019/file/ddf354219aac374f1d40b7e760ee5bb7-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/ddf354219aac374f1d40b7e760ee5bb7-Paper.pdf).
- 500 Nikita Balabin, Daria Voronkova, Ilya Trofimov, Evgeny Burnaev, and Serguei Barannikov. Disen-
 501 tanglement learning via topology. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian
 502 Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st
 503 International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learn-
 504 ing Research*, pp. 2474–2504. PMLR, 21–27 Jul 2024. URL [https://proceedings.mlr.
 505 press/v235/balabin24a.html](https://proceedings.mlr.press/v235/balabin24a.html).
- 506 Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization
 507 for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- 508 Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron
 509 Courville, and R Devon Hjelm. Mine: Mutual information neural estimation, 2021. URL
 510 <https://arxiv.org/abs/1801.04062>.
- 511 A J Bell and T J Sejnowski. An information-maximization approach to blind separation and blind
 512 deconvolution. *Neural Comput.*, 7(6):1129–1159, November 1995.
- 513 Vanessa Böhm and Uroš Seljak. Probabilistic autoencoder. *Transactions on Machine Learning
 514 Research*, 2022. doi: 10.48550/ARXIV.2006.05479. URL [https://openreview.net/
 515 forum?id=AEoYjvjKVA](https://openreview.net/forum?id=AEoYjvjKVA).
- 516 Johann Brehmer and Kyle Cranmer. Flows for simultaneous manifold learning and density esti-
 517 mation. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances
 518 in Neural Information Processing Systems*, volume 33, pp. 442–453. Curran Associates, Inc.,
 519 2020. URL [https://proceedings.neurips.cc/paper_files/paper/2020/
 520 file/051928341be67dcb03f0e04104d9047-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/051928341be67dcb03f0e04104d9047-Paper.pdf).
- 521 Bradley CA Brown, Anthony L. Caterini, Brendan Leigh Ross, Jesse C Cresswell, and Gabriel
 522 Loaiza-Ganem. Verifying the union of manifolds hypothesis for image data. In *The Eleventh
 523 International Conference on Learning Representations*, 2023. URL [https://openreview.
 524 net/forum?id=Rvee9CAX4fi](https://openreview.net/forum?id=Rvee9CAX4fi).
- 525 Ivan Butakov, Alexander Tolmachev, Sofia Malanchuk, A. M. Neopryatnaya, and Alexey Frolov.
 526 Mutual information estimation via normalizing flows. *ArXiv*, abs/2403.02187, 2024a. URL
 527 <https://api.semanticscholar.org/CorpusID:268247846>.
- 528 Ivan Butakov, Alexander Tolmachev, Sofia Malanchuk, Anna Neopryatnaya, Alexey Frolov, and
 529 Kirill Andreev. Information bottleneck analysis of deep neural networks via lossy compression.
 530 In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=huGECz8dPp>.
- 531 Scott Chen and Ramesh Gopinath. Gaussianization. In T. Leen, T. Dietterich, and
 532 V. Tresp (eds.), *Advances in Neural Information Processing Systems*, volume 13. MIT Press,
 533 2000. URL [https://proceedings.neurips.cc/paper_files/paper/2000/
 534 file/3c947bc2f7ff007b86a9428b74654de5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2000/file/3c947bc2f7ff007b86a9428b74654de5-Paper.pdf).

- 540 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework
541 for contrastive learning of visual representations. In Hal Daumé III and Aarti Singh (eds.),
542 *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Pro-*
543 *ceedings of Machine Learning Research*, pp. 1597–1607. PMLR, 13–18 Jul 2020a. URL
544 <https://proceedings.mlr.press/v119/chen20j.html>.
- 545 Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-
546 supervised models are strong semi-supervised learners. In *Proceedings of the 34th International*
547 *Conference on Neural Information Processing Systems*, NIPS ’20, Red Hook, NY, USA, 2020b.
548 Curran Associates Inc. ISBN 9781713829546.
- 549 Kai Lai Chung. *A course in probability theory*. Academic Press, 3rd ed edition, 2001. ISBN
550 0121741516; 9780121741518; 9780080522982; 008052298X.
- 551 Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecom-*
552 *munications and Signal Processing)*. Wiley-Interscience, USA, 2006.
- 553 Victor Guilherme Turrise da Costa, Enrico Fini, Moin Nabi, Nicu Sebe, and Elisa Ricci. solo-
554 learn: A library of self-supervised methods for visual representation learning. *Journal of Machine*
555 *Learning Research*, 23(56):1–6, 2022. URL [http://jmlr.org/papers/v23/21-1155.](http://jmlr.org/papers/v23/21-1155.html)
556 [html](http://jmlr.org/papers/v23/21-1155.html).
- 557 F. N. David and N. L. Johnson. The probability integral transformation when parameters are
558 estimated from the sample. *Biometrika*, 35(1/2):182–190, 1948. ISSN 00063444. URL
559 <http://www.jstor.org/stable/2332638>.
- 560 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep
561 bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and
562 Tamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of*
563 *the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*
564 *2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pp. 4171–
565 4186. Association for Computational Linguistics, 2019. doi: 10.18653/V1/N19-1423. URL
566 <https://doi.org/10.18653/v1/n19-1423>.
- 567 M. D. Donsker and S. R.S. Varadhan. Asymptotic evaluation of certain markov process expectations
568 for large time. iv. *Communications on Pure and Applied Mathematics*, 36(2):183–212, March
569 1983. ISSN 0010-3640. doi: 10.1002/cpa.3160360204.
- 570 B. Duong and T. Nguyen. Conditional independence testing via latent representation learning. In
571 *2022 IEEE International Conference on Data Mining (ICDM)*, pp. 121–130, Los Alamitos, CA,
572 USA, dec 2022. IEEE Computer Society. doi: 10.1109/ICDM54844.2022.00022. URL [https:](https://doi.ieeecomputersociety.org/10.1109/ICDM54844.2022.00022)
573 [//doi.ieeecomputersociety.org/10.1109/ICDM54844.2022.00022](https://doi.ieeecomputersociety.org/10.1109/ICDM54844.2022.00022).
- 574 R. B. D’Agostino. An omnibus test of normality for moderate and large sample size. *Biometrika*,
575 58(1):341–348, 1971.
- 576 R. B. D’Agostino and E. S. Pearson. Tests for departure from normality. *Biometrika*, 60(1):613–622,
577 1973.
- 578 Linus Ericsson, Henry Gouk, Chen Change Loy, and Timothy M. Hospedales. Self-supervised rep-
579 resentation learning: Introduction, advances, and challenges. *IEEE Signal Processing Magazine*,
580 39(3):42–62, 2022. doi: 10.1109/MSP.2021.3134634.
- 581 Marco Federici, Anjan Dutta, Patrick Forré, Nate Kushman, and Zeynep Akata. Learning robust
582 representations via multi-view information bottleneck. In *International Conference on Learning*
583 *Representations*, 2020. URL <https://openreview.net/forum?id=B1xwcyHFDr>.
- 584 Elizabeth González-Estrada, José Villaseñor, and Rocío Acosta-Pech. Shapiro-wilk test for
585 multivariate skew-normality. *Computational Statistics*, 37:1–17, 09 2022. doi: 10.1007/
586 s00180-021-01188-y.
- 587 Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
588 Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. URL [https:](https://arxiv.org/abs/1406.2661)
589 [//arxiv.org/abs/1406.2661](https://arxiv.org/abs/1406.2661).

- 594 Wes Gurnee and Max Tegmark. Language models represent space and time. In *The Twelfth Interna-*
595 *tional Conference on Learning Representations*, 2024. URL [https://openreview.net/](https://openreview.net/forum?id=jE8xbmvFin)
596 [forum?id=jE8xbmvFin](https://openreview.net/forum?id=jE8xbmvFin).
597
- 598 Robert M. Haralick, K. Shanmugam, and Its’Hak Dinstein. Textural features for image classification.
599 *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-3(6):610–621, 1973. doi: 10.1109/
600 TSMC.1973.4309314.
- 601 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
602 *nition*, 2015. URL <https://arxiv.org/abs/1512.03385>.
603
- 604 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
605 *unsupervised visual representation learning*. In *2020 IEEE/CVF Conference on Computer Vision*
606 *and Pattern Recognition (CVPR)*, pp. 9726–9735, 2020. doi: 10.1109/CVPR42600.2020.00975.
- 607 Norbert Henze and Bernd Zirkler. A class of invariant consistent tests for multivariate normality.
608 *Communications in Statistics-theory and Methods*, 19:3595–3617, 1990. URL [https://api.](https://api.semanticscholar.org/CorpusID:120328121)
609 [semanticscholar.org/CorpusID:120328121](https://api.semanticscholar.org/CorpusID:120328121).
610
- 611 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochre-
612 *iter*. Gans trained by a two time-scale update rule converge to a local nash equilibrium.
613 In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and
614 R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran
615 Associates, Inc., 2017. URL [https://proceedings.neurips.cc/paper_files/](https://proceedings.neurips.cc/paper_files/paper/2017/file/8ald694707eb0fefe65871369074926d-Paper.pdf)
616 [paper/2017/file/8ald694707eb0fefe65871369074926d-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/8ald694707eb0fefe65871369074926d-Paper.pdf).
- 617 Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick,
618 Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a
619 *constrained variational framework*. In *International Conference on Learning Representations*,
620 2017. URL <https://openreview.net/forum?id=Sy2fzU9gl>.
- 621 R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam
622 Trischler, and Yoshua Bengio. Learning deep representations by mutual information estima-
623 *tion and maximization*. In *International Conference on Learning Representations*, 2019. URL
624 <https://openreview.net/forum?id=Bklr3j0cKX>.
625
- 626 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on*
627 *Deep Generative Models and Downstream Applications*, 2021. URL [https://openreview.](https://openreview.net/forum?id=qw8AKxfYbI)
628 [net/forum?id=qw8AKxfYbI](https://openreview.net/forum?id=qw8AKxfYbI).
- 629 Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron Courville. Neural autoregressive
630 *flows*. In *International conference on machine learning*, pp. 2078–2087. PMLR, 2018.
631
- 632 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. Position: The platonic represen-
633 *tation hypothesis*. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria
634 Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International*
635 *Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*,
636 pp. 20617–20642. PMLR, 21–27 Jul 2024. URL [https://proceedings.mlr.press/](https://proceedings.mlr.press/v235/huh24a.html)
637 [v235/huh24a.html](https://proceedings.mlr.press/v235/huh24a.html).
- 638 Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by
639 *reducing internal covariate shift*, 2015. URL <https://arxiv.org/abs/1502.03167>.
640
- 641 Priyank Jaini, Kira A. Selby, and Yaoliang Yu. Sum-of-squares polynomial flow. In Kamalika
642 Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on*
643 *Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3009–3018.
644 PMLR, 09–15 Jun 2019.
- 645 Amor Keziou. Dual representation of ϕ -divergences and applications. *Comptes Ren-*
646 *endus Mathematique*, 336(10):857–862, 2003. ISSN 1631-073X. doi: [https://doi.org/10.](https://doi.org/10.1016/S1631-073X(03)00215-2)
647 [1016/S1631-073X\(03\)00215-2](https://doi.org/10.1016/S1631-073X(03)00215-2). URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S1631073X03002152)
[article/pii/S1631073X03002152](https://www.sciencedirect.com/science/article/pii/S1631073X03002152).

- 648 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017.
649
- 650 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann
651 LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB,
652 Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL [http://arxiv.org/
653 abs/1312.6114](http://arxiv.org/abs/1312.6114).
- 654 Ivan Kobyzev, Simon Prince, and Marcus A. Brubaker. Normalizing flows: An introduction
655 and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelli-
656 gence*, 43:3964–3979, 2020. URL [https://api.semanticscholar.org/CorpusID:
657 208910764](https://api.semanticscholar.org/CorpusID:208910764).
- 658
- 659 Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
660
- 661 Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Au-
662 toencoding beyond pixels using a learned similarity metric. In Maria Florina Balcan and Kil-
663 ian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learn-
664 ing*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1558–1566, New York, New
665 York, USA, 20–22 Jun 2016. PMLR. URL [https://proceedings.mlr.press/v48/
666 larsen16.html](https://proceedings.mlr.press/v48/larsen16.html).
- 667 Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online].
668 Available: http://yann.lecun.com/exdb/mnist*, 2, 2010.
669
- 670 R. Linsker. Self-organization in a perceptual network. *Computer*, 21(3):105–117, 1988. doi: 10.
671 1109/2.36.
- 672
- 673 Alireza Makhzani and Brendan Frey. k-sparse autoencoders. In Yoshua Bengio and Yann Le-
674 Cun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB,
675 Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL [https://arxiv.
676 org/abs/1312.5663](https://arxiv.org/abs/1312.5663).
- 677 Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, and Ian Goodfellow. Adversarial autoencoders.
678 In *International Conference on Learning Representations*, 2016. URL [http://arxiv.org/
679 abs/1511.05644](http://arxiv.org/abs/1511.05644).
- 680
- 681 David McAllester and Karl Stratos. Formal limitations on the measurement of mutual informa-
682 tion. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third Interna-
683 tional Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine
684 Learning Research*, pp. 875–884. PMLR, 26–28 Aug 2020. URL [https://proceedings.
685 mlr.press/v108/mcallester20a.html](https://proceedings.mlr.press/v108/mcallester20a.html).
- 686 Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. Umap: Uniform manifold
687 approximation and projection. *Journal of Open Source Software*, 3(29):861, 2018. doi: 10.
688 21105/joss.00861. URL <https://doi.org/10.21105/joss.00861>.
- 689
- 690 Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word represen-
691 tations in vector space, 2013. URL <https://arxiv.org/abs/1301.3781>.
- 692
- 693 Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets, 2014. URL [https:
694 //arxiv.org/abs/1411.1784](https://arxiv.org/abs/1411.1784).
- 695
- 696 Michael Moor, Max Horn, Bastian Rieck, and Karsten Borgwardt. Topological autoencoders. In
697 Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Ma-
698 chine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7045–7054.
699 PMLR, 13–18 Jul 2020. URL [https://proceedings.mlr.press/v119/moor20a.
700 html](https://proceedings.mlr.press/v119/moor20a.html).
- 701 Andrew Ng. Sparse autoencoder, 2011. URL [http://www.stanford.edu/class/
cs294a/sparseAutoencoder.pdf](http://www.stanford.edu/class/cs294a/sparseAutoencoder.pdf).

- 702 XuanLong Nguyen, Martin J. Wainwright, and Michael I. Jordan. Estimating divergence functionals
703 and the likelihood ratio by convex risk minimization. *IEEE Trans. Inf. Theor.*, 56(11):5847–5861,
704 nov 2010. ISSN 0018-9448. doi: 10.1109/TIT.2010.2068870. URL <https://doi.org/10.1109/TIT.2010.2068870>.
705
- 706 George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lak-
707 shminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Ma-
708 chine Learning Research*, 22(57):1–64, 2021. URL [http://jmlr.org/papers/v22/
709 19-1028.html](http://jmlr.org/papers/v22/19-1028.html).
710
- 711 Y. Polyanskiy and Y. Wu. *Information Theory: From Coding to Learning*. Cambridge Univer-
712 sity Press, 2024. ISBN 9781108832908. URL [https://books.google.ru/books?id=
713 CySo0AEACAAJ](https://books.google.ru/books?id=CySo0AEACAAJ).
- 714 Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language under-
715 standing by generative pre-training, 2018.
716
- 717 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agar-
718 wal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya
719 Sutskever. Learning transferable visual models from natural language supervision, 2021. URL
720 <https://arxiv.org/abs/2103.00020>.
- 721 R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image syn-
722 thesis with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision
723 and Pattern Recognition (CVPR)*, pp. 10674–10685, Los Alamitos, CA, USA, jun 2022.
724 IEEE Computer Society. doi: 10.1109/CVPR52688.2022.01042. URL [https://doi.
725 ieecomputersociety.org/10.1109/CVPR52688.2022.01042](https://doi.ieeeecomputersociety.org/10.1109/CVPR52688.2022.01042).
726
- 727 Peter J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of clus-
728 ter analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 11 1987. doi:
729 10.1016/0377-0427(87)90125-7.
- 730 Peter Sorrenson, Felix Draxler, Armand Rousselot, Sander Hummerich, Lea Zimmermann, and
731 Ullrich Koethe. Lifting architectural constraints of injective flows. In *The Twelfth International
732 Conference on Learning Representations*, 2024. URL [https://openreview.net/forum?
733 id=kBNix4Biq4](https://openreview.net/forum?id=kBNix4Biq4).
- 734 M. Spivak. *Calculus On Manifolds: A Modern Approach To Classical Theorems Of Advanced
735 Calculus*. Avalon Publishing, 1965. ISBN 9780805390216.
736
- 737 Michael E. Tipping and Christopher M. Bishop. Mixtures of Probabilistic Principal Component
738 Analyzers. *Neural Computation*, 11(2):443–482, 02 1999. ISSN 0899-7667. doi: 10.1162/
739 089976699300016728. URL <https://doi.org/10.1162/089976699300016728>.
- 740 Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck
741 method. In *Proc. of the 37-th Annual Allerton Conference on Communication, Control
742 and Computing*, pp. 368–377, 1999. URL [/brokenurl#citeseeer.nj.nec.com/
743 tishby99information.html](/brokenurl#citeseeer.nj.nec.com/tishby99information.html).
744
- 745 Antonio Trujillo-Ortiz, Rafael Hernandez-Walls, and K. Barba-Rojo. Hzmvttest. henze-zirkler’s
746 multivariate normality test., 12 2007.
747
- 748 Michael Tschannen, Josip Djolonga, Paul K. Rubenstein, Sylvain Gelly, and Mario Lucic. On mu-
749 tual information maximization for representation learning. In *International Conference on Learn-
750 ing Representations*, 2020. URL <https://openreview.net/forum?id=rkxoh24FPH>.
- 751 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predic-
752 tive coding, 2019. URL <https://arxiv.org/abs/1807.03748>.
753
- 754 Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Ma-
755 chine Learning Research*, 9(86):2579–2605, 2008. URL [http://jmlr.org/papers/v9/
vandermaaten08a.html](http://jmlr.org/papers/v9/vandermaaten08a.html).

756 Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R Devon
757 Hjelm. Deep graph infomax. In *International Conference on Learning Representations*, 2019.
758 URL <https://openreview.net/forum?id=rklz9iAcKQ>.

759 Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural
760 image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern
761 recognition*, pp. 3156–3164, 2015.

762 Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks, 2017.
763 URL <https://arxiv.org/abs/1708.03888>.

764 Xuechu Yu. Leveraging superfluous information in contrastive representation learning, 2024. URL
765 <https://arxiv.org/abs/2408.10292>.

766 Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised
767 learning via redundancy reduction. In *International conference on machine learning*, pp. 12310–
768 12320. PMLR, 2021.

772 A COMPLETE PROOFS

773 **Theorem 2.1** (Theorem 8.6.5 in Cover & Thomas (2006)). *Let X be a d -dimensional absolutely
774 continuous random vector with probability density function p , mean m and covariance matrix Σ .
775 Then*

$$776 h(X) = h(\mathcal{N}(m, \Sigma)) - D_{\text{KL}}(p \parallel \mathcal{N}(m, \Sigma)), \quad h(\mathcal{N}(m, \Sigma)) = \frac{1}{2} \log((2\pi e)^d \det \Sigma),$$

777 where $\mathcal{N}(m, \Sigma)$ is a Gaussian distribution of mean m and covariance matrix Σ .

778 *Proof of Theorem 2.1.* As for any $m \in \mathbb{R}^d$ it holds $h(X - m) = h(X)$, let us consider a centered
779 random vector X . Denoting the probability density function of $\mathcal{N}(0, \Sigma)$ by ϕ_Σ , we have

$$780 D_{\text{KL}}(p \parallel \mathcal{N}(0, \Sigma)) = \int_{\mathbb{R}^d} p(x) \log \frac{p(x)}{\phi_\Sigma(x)} dx = -h(X) - \int_{\mathbb{R}^d} p(x) \log \phi_\Sigma(x) dx.$$

781 Now, consider the second term:

$$782 \int_{\mathbb{R}^d} p(x) \log \phi_\Sigma(x) dx = \text{const} + \frac{1}{2} \mathbb{E}_X X^T \Sigma^{-1} X =$$

$$783 = \text{const} + \frac{1}{2} \text{Tr}(\Sigma^{-1} \mathbb{E}_X [X X^T]) = \text{const} + \frac{1}{2} \text{Tr}(\Sigma^{-1} \Sigma) = \text{const} + \frac{d}{2} =$$

$$784 = \text{const} + \frac{1}{2} \mathbb{E}_{\mathcal{N}(0, \Sigma)} X^T \Sigma^{-1} X = \int_{\mathbb{R}^d} \phi_\Sigma(x) \log \phi_\Sigma(x) dx.$$

785 Here, in the second line, the cyclic property of the trace is used. \square

786 **Theorem 2.2.** *Let X be an absolutely continuous random vector with probability density function
787 p and $\text{supp } X \subseteq S$, where S has finite and non-zero Lebesgue measure $\mu(S)$. Then*

$$788 h(X) = h(\text{U}(S)) - D_{\text{KL}}(p \parallel \text{U}(S)), \quad h(\text{U}(S)) = \log \mu(S),$$

789 where $\text{U}(S)$ is a uniform distribution on S .

790 *Proof of Theorem 2.2.*

$$791 D_{\text{KL}}(p \parallel \text{U}(S)) = \int_S p(x) \log \frac{p(x)}{1/\mu(S)} dx =$$

$$792 = -h(X) + \int_S p(x) \log \mu(S) dx = -h(X) + \log \mu(S) = -h(X) + h(\text{U}(S)).$$

793 \square

Lemma 2.4. *Let X and Z be random vectors of the same dimensionality, independent under the conditioning vector Y . Then*

$$h(X + Z | Y) = h(Z | Y) + I(X; X + Z | Y) \geq h(Z | Y),$$

with equality if and only if there exists a measurable function g such that $X = g(Y)$.³

Proof of Lemma 2.4. By the definition of mutual information between Z and $X + Z$ conditioned on Y , we have

$$I(X; X + Z | Y) = h(X + Z | Y) - h(X + Z | X, Y),$$

where $h(X + Z | X, Y) = h(Z | Y)$ due to the independence of X and Z given Y .

Now, let us prove that $I(X; X + Z | Y) = 0$ if and only if $\exists g : X = g(Y)$. Firstly, if $X = g(Y)$,

$$I(X; X + Z | Y) = I(g(Y); g(Y) + Z | Y) = I(0; Z | Y) = I(0; Z) = 0$$

Next, consider $I(X; X + Z | Y) = 0$. Thus, X is conditionally independent both of $X + Z$ and Z , which allows for the multiplicative property of characteristic functions to be used:

$$\begin{aligned} \mathbb{E} \left(e^{i\langle t, Z \rangle} \mid Y \right) &= \mathbb{E} \left(e^{i\langle t, -X + (Z+X) \rangle} \mid Y \right) = \mathbb{E} \left(e^{-i\langle t, X \rangle} \mid Y \right) \mathbb{E} \left(e^{i\langle t, Z \rangle} \mid Y \right) \mathbb{E} \left(e^{i\langle t, X \rangle} \mid Y \right) = \\ &= \left[\mathbb{E} \left(e^{i\langle t, X \rangle} \mid Y \right) \right]^* \mathbb{E} \left(e^{i\langle t, X \rangle} \mid Y \right) \mathbb{E} \left(e^{i\langle t, Z \rangle} \mid Y \right) = \left| \mathbb{E} \left(e^{i\langle t, X \rangle} \mid Y \right) \right|^2 \mathbb{E} \left(e^{i\langle t, Z \rangle} \mid Y \right) \end{aligned}$$

Note that there exists $\varepsilon(Y) > 0$ such that $\mathbb{E} \left(e^{i\langle t, Z \rangle} \mid Y \right) \neq 0$ for $\|t\| < \varepsilon(Y)$ (non-vanishing property). Consequently, $\left| \mathbb{E} \left(e^{i\langle t, X \rangle} \mid Y \right) \right| = 1$ holds for $\|t\| < \varepsilon(Y)$. This implies that the conditional distribution of X given Y is a δ -distribution (Theorem 6.4.7 in (Chung, 2001)), leading to $X = \mathbb{E}(X | Y)$, where $g(Y) \triangleq \mathbb{E}(X | Y)$ is a $\sigma(Y)$ -measurable function. \square

Lemma 3.1. *Consider the following Markov chain of absolutely continuous random vectors:*

$$f(X) + Z \longrightarrow X \longrightarrow X' \longrightarrow f(X'),$$

with Z being independent of (X, X') . Then

$$I(f(X'); f(X) + Z) = h(f(X) + Z) - h(Z) - I(f(X) + Z; f(X) | f(X')). \quad (5)$$

Proof of Lemma 3.1. From the definition of mutual information, we have

$$I(f(X'); f(X) + Z) = h(f(X) + Z) - h(f(X) + Z | f(X')).$$

We apply Lemma 2.4 to rewrite the second term

$$\begin{aligned} h(f(X) + Z | f(X')) &= h(Z | f(X')) + I(f(X); f(X) + Z | f(X')) \\ &= h(Z) + I(f(X); f(X) + Z | f(X')), \end{aligned}$$

where the independence of Z and $f(X')$ is used. \square

Lemma 3.3. *Under the conditions of Lemma 3.1, let $\mathbb{P}(X = X' | X) \geq \alpha > 0$. Then, $I(f(X) + Z; f(X) | f(X')) = 0$ precisely when f is weakly invariant to $X \rightarrow X'$.*

Proof of Lemma 3.3. According to Lemma 2.4, $I(f(X) + Z; f(X) | f(X')) = 0$ if and only if $f(X) = g(f(X'))$ for some g . Thus, weak invariance implies $I(f(X) + Z; f(X) | f(X')) = 0$.

On the other hand, if $f(X) = g(f(X'))$,

$$\mathbb{P}(f(X) \in (g^{-1} \circ f)(X) | X) \geq \mathbb{P}(f(X) = f(X') | X) \geq \mathbb{P}(X = X' | X) \geq \alpha > 0.$$

Note that the predicate $P(X) \triangleq "f(X) \in (g^{-1} \circ f)(X)"$ is not random when conditioned on X . Thus, $\mathbb{P}(P(X) | X) = \phi(X)$, where ϕ is a function taking values in $\{0, 1\}$. As $\phi(X) \geq \alpha > 0$, $\phi(X) = 1$. This implies $g(f(X)) = f(X)$ almost surely. \square

³Hereinafter, when comparing random variables, we mean equality or inequality "almost sure".

Theorem 3.4 (Gaussian distribution matching). *Let the conditions of Lemma 3.3 be satisfied. Assume $Z \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, $\mathbb{E} f(X) = 0$ and $\text{Var}(f(X)_i) = 1$ for all $i \in \{1, \dots, d\}$. Then, the mutual information $I(f(X'); f(X) + Z)$ can be upper bounded as follows*

$$I(f(X'); f(X) + Z) \leq \frac{d}{2} \log \left(1 + \frac{1}{\sigma^2} \right), \quad (6)$$

with the equality holding exactly when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$. Moreover,

$$D_{\text{KL}}(f(X) \parallel \mathcal{N}(0, \mathbf{I})) \leq I(Z; f(X) + Z) - I(f(X'); f(X) + Z) - d \log \sigma.$$

Proof of Theorem 3.4. Applying the result from Lemma B.1 gives us:

$$D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) \leq \frac{d}{2} \log \left(1 + \frac{1}{\sigma^2} \right) - I(f(X'); f(X) + Z).$$

Since KL-divergence is non-negative, we obtain the desired bound. Furthermore, equality holds exactly when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$, as discussed in the proof of Lemma B.1.

The second inequality involving $D_{\text{KL}}(f(X) \parallel \mathcal{N}(0, \mathbf{I}))$, follows from Corollary B.3. \square

Theorem 3.5 (Uniform distribution matching). *Under the conditions of Lemma 3.1, let $Z \sim \text{U}([-\varepsilon; \varepsilon]^d)$ and $\text{supp } f(X) \subseteq [0; 1]^d$. Then, the mutual information $I(f(X'); f(X) + Z)$ can be upper bounded as follows*

$$I(f(X'); f(X) + Z) \leq d \log \left(1 + \frac{1}{2\varepsilon} \right). \quad (7)$$

with the equality if and only if $1/\varepsilon \in \mathbb{N}$, f is weakly invariant, and $f(X) \sim \text{U}(A)$, where the set $A = \{0, 2\varepsilon, 4\varepsilon, \dots, 1\}$ contains $(1/(2\varepsilon) + 1)$ elements.

Moreover,

$$D_{\text{KL}}(f(X) \parallel \text{U}([0; 1]^d)) \leq I(Z; f(X) + Z) - I(f(X'); f(X) + Z) - d \log(2\varepsilon).$$

Proof of Theorem 3.5. Proceeding similarly to the Gaussian case, we apply Lemma B.4, yielding

$$D_{\text{KL}}(f(X) + Z \parallel \text{U}([-\varepsilon; 1 + \varepsilon]^d)) \leq d \log \left(1 + \frac{1}{2\varepsilon} \right) - I(f(X'); f(X) + Z),$$

where the left-hand side is non-negative, so we obtain the claimed inequality. The conditions for equality are also established through Lemma B.4.

To prove the second upper-bound, concerning $D_{\text{KL}}(f(X) \parallel \text{U}([0; 1]^d))$, it suffices to use the result from Corollary B.6. \square

Theorem 4.1 (Dual form of Gaussian distribution matching). *Under the conditions of Theorem 3.4,*

$$I(f(X'); f(X) + Z) \geq \mathbb{E}_{\mathbb{P}^+} \left[T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^* \right] - \log \mathbb{E}_{\mathbb{P}^-} \left[\exp \left(T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^* \right) \right],$$

$$T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^*(x, y) = \frac{\|y\|^2}{2(1 + \sigma^2)} - \frac{\|y - x\|^2}{2\sigma^2} = \frac{1}{\sigma^2} \left(\langle x, y \rangle - \frac{\|x\|^2 + \|y\|^2}{2(1 + \sigma^2)} \right),$$

with the equality holding precisely when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$.

Proof of Theorem 4.1. By Remark 2.5 we know that the equality holds if and only if $T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^*(x, y) = \text{PMI}_{f(X'), f(X)+Z}(x, y) + \text{const.}$

Now, for independent $Y \sim \mathcal{N}(0, \mathbf{I})$ and $Z \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$

$$\begin{aligned} \text{PMI}_{Y, Y+Z}(x, y) &= \log p_{Y+Z|Y}(y | x) - \log p_{Y+Z}(y) = \\ &= \frac{\|y\|^2}{2(1 + \sigma^2)} - \frac{\|y - x\|^2}{2\sigma^2} + \frac{d}{2} \log \left(1 + \frac{1}{\sigma^2} \right) = \\ &= T_{\mathcal{N}(0, \sigma^2 \mathbf{I})}^*(x, y) + I(Y; Y + Z). \end{aligned}$$

Thus, the equality holds if and only if $I(f(X'); f(X) + Z) = \frac{d}{2} \log(1 + 1/\sigma^2)$, which, in turn, holds precisely when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$ (see Theorem 3.4). \square

B SUPPLEMENTARY RESULTS ON DISTRIBUTION MATCHING

In this section, we explore the connection between the proposed infomax objective and the problem of distribution matching, formulated in terms of the Kullback-Leibler divergence.

GAUSSIAN DISTRIBUTION MATCHING

Lemma B.1. *Assume the conditions of Theorem 3.4 are satisfied, then for Gaussian distribution matching, we have*

$$D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) \leq \frac{d}{2} \log \left(1 + \frac{1}{\sigma^2} \right) - I(f(X'); f(X) + Z),$$

with equality holding exactly when f is weakly invariant and $f(X) \sim \mathcal{N}(0, \mathbf{I})$.

Proof. From Lemma 3.1 we obtain

$$I(f(X'); f(X) + Z) = h(f(X) + Z) - h(\mathcal{N}(0, \sigma^2\mathbf{I})) - I(f(X) + Z; f(X) \mid f(X')).$$

Using Theorem 2.1, we can rewrite the first term, which yields

$$\begin{aligned} I(f(X'); f(X) + Z) &= h(\mathcal{N}(m, \Sigma)) - h(\mathcal{N}(0, \sigma^2\mathbf{I})) \\ &\quad - D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(m, \Sigma)) - I(f(X) + Z; f(X) \mid f(X')), \end{aligned}$$

where m and Σ are the mean and covariance matrix of $f(X) + Z$.

To bound the KL-divergence, note that the conditional mutual information is non-negative:

$$D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(m, \Sigma)) \leq h(\mathcal{N}(m, \Sigma)) - h(\mathcal{N}(0, \sigma^2\mathbf{I})) - I(f(X'); f(X) + Z).$$

Equality holds exactly when $I(f(X) + Z; f(X) \mid f(X')) = 0$, which is equivalent to f being weakly invariant (see Lemma 3.3).

Next, we estimate the difference between the entropies by observing that

$$h(\mathcal{N}(m, \Sigma)) \leq \sum_{i=1}^d h(\mathcal{N}(m_i, \text{Var}(f(X)_i) + \sigma^2)) = d \cdot h(\mathcal{N}(0, 1 + \sigma^2)),$$

with the equality holding if and only if Σ is diagonal, which implies $\Sigma = \mathbf{I}$ since $\text{Var}(f(X)_i) = 1$ for all $i \in \{1, \dots, d\}$. Finally,

$$d \cdot h(\mathcal{N}(0, 1 + \sigma^2)) - h(\mathcal{N}(0, \sigma^2\mathbf{I})) = \frac{d}{2} [\log(1 + \sigma^2) - \log \sigma^2] = \frac{d}{2} \log \left(1 + \frac{1}{\sigma^2} \right),$$

which proves the claimed inequality. \square

Lemma B.2. *Under the conditions of Theorem 3.4, the following holds:*

$$\begin{aligned} D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) &\leq D_{\text{KL}}(f(X) \parallel \mathcal{N}(0, \mathbf{I})) = \\ &= D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) + I(Z; f(X) + Z) - \frac{d}{2} \log(1 + \sigma^2). \end{aligned}$$

Proof. The left-hand inequality follows directly from the data processing inequality for the Kullback-Leibler divergence (Theorem 2.17 in (Polyanskiy & Wu, 2024)).

To establish the right-hand side, we first apply Theorem 2.1, and then use the independence of $f(X)$ and Z , yielding $h(f(X)) = h(f(X) + Z \mid Z)$. Thus, we have

$$D_{\text{KL}}(f(X) \parallel \mathcal{N}(0, \mathbf{I})) = h(\mathcal{N}(0, \mathbf{I})) - h(f(X) + Z \mid Z).$$

Next, by using the definition of mutual information and again applying Theorem 2.1, one can write

$$\begin{aligned} h(f(X) + Z \mid Z) &= h(f(X) + Z) - I(f(X) + Z; Z) = \\ &= h(\mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) - D_{\text{KL}}(f(X) + Z \parallel \mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) - I(f(X) + Z; Z). \end{aligned}$$

Finally, substituting this into the expansion for $D_{\text{KL}}(f(X) \parallel \mathcal{N}(0, \mathbf{I}))$, and noting that $h(\mathcal{N}(0, \mathbf{I})) - h(\mathcal{N}(0, (1 + \sigma^2)\mathbf{I})) = -\frac{d}{2} \log(1 + \sigma^2)$ completes the proof. \square

972 **Corollary B.3.** *In the Gaussian distribution matching setup (Theorem 3.4), we have*

$$973 \quad D_{\text{KL}}(f(X) \parallel \mathcal{N}(0, \mathbf{I})) \leq I(Z; f(X) + Z) - I(f(X'); f(X) + Z) - d \log \sigma.$$

975 *Proof.* This follows directly from combining Lemma B.1 and Lemma B.2. □

977 UNIFORM DISTRIBUTION MATCHING

979 **Lemma B.4.** *Let the conditions of Theorem 3.5 hold, then the following bound applies for uniform*
 980 *distribution matching:*

$$981 \quad D_{\text{KL}}(f(X) + Z \parallel \mathbf{U}([- \varepsilon; 1 + \varepsilon]^d)) \leq d \log \left(1 + \frac{1}{2\varepsilon} \right) - I(f(X'); f(X) + Z),$$

982 *with the equality if and only if $1/\varepsilon \in \mathbb{N}$, f is weakly invariant, and $f(X) \sim \mathbf{U}(A)$, where the set*
 983 *$A = \{0, 2\varepsilon, 4\varepsilon, \dots, 1\}$ contains $(1/(2\varepsilon) + 1)$ elements.*

984 *Proof.* Similarly to the proof of Lemma B.1, we use the decomposition of the infomax objective
 985 from Lemma 3.1. Afterward, we apply Theorem 2.2 to the term $h(f(X) + Z)$:

$$986 \quad I(f(X'); f(X) + Z) = h(\mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) - h(\mathbf{U}([- \varepsilon; \varepsilon]^d)) \\ 987 \quad - D_{\text{KL}}(f(X) + Z \parallel \mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) - I(f(X) + Z; f(X) \mid f(X')).$$

988 Therefore, the KL-divergence can be bounded as:

$$989 \quad D_{\text{KL}}(f(X) + Z \parallel \mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) \leq h(\mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) - h(\mathbf{U}([- \varepsilon; \varepsilon]^d)) - I(f(X'); f(X) + Z) \\ 990 \quad = d \log \left(1 + \frac{1}{2\varepsilon} \right) - I(f(X'); f(X) + Z),$$

991 with equality achieved if and only if the mutual information between $f(X) + Z$ and $f(X)$ condi-
 992 tioned on $f(X')$ is zero, which occurs precisely when f is weakly invariant (see Lemma 3.3).

993 Next, we show when the equality holds. The probability density function of the sum of independent
 994 random vectors $f(X)$ and Z is given by the convolution:

$$995 \quad \int_{\mathbb{R}^d} p_{f(X)}(x) p_Z(z-x) dx = \prod_{i=1}^d \int_{z_i-\varepsilon}^{z_i+\varepsilon} \left(\frac{1}{1/(2\varepsilon) + 1} \sum_{a \in A} \delta(z_i - a) \right) \frac{dx_i}{2\varepsilon} = \frac{1}{(1 + 2\varepsilon)^d} \int_{[-\varepsilon; 1+\varepsilon]^d} dx.$$

996 Here, we used the independence of the components, with Z_i uniformly distributed on $[-\varepsilon, \varepsilon]$ and
 997 $f(X)_i$ uniformly distributed over the discrete set $A = \{0, 2\varepsilon, 4\varepsilon, \dots, 1\}$.

998 Therefore, $(f(X) + Z) \sim \mathbf{U}([- \varepsilon; 1 + \varepsilon]^d)$. Given this, one can calculate the mutual information
 999 explicitly:

$$1000 \quad I(f(X); f(X) + Z) = h(f(X) + Z) - h(Z) = d \log \left(1 + \frac{1}{2\varepsilon} \right).$$

1001 This concludes the proof. □

1002 **Lemma B.5.** *Under the conditions of Theorem 3.5, the following holds:*

$$1003 \quad D_{\text{KL}}(f(X) \parallel \mathbf{U}([0; 1]^d)) = D_{\text{KL}}(f(X) + Z \parallel \mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) + I(Z; f(X) + Z) - d \log(1 + 2\varepsilon).$$

1004 *Proof.* One can build on the reasoning from Lemma B.2 by applying Theorem 2.2 to express the
 1005 KL-divergence between $f(X)$ and $\mathbf{U}([0; 1]^d)$ as follows:

$$1006 \quad D_{\text{KL}}(f(X) \parallel \mathbf{U}([0; 1]^d)) = h(\mathbf{U}([0; 1]^d)) - h(f(X) + Z \mid Z) = -h(f(X) + Z \mid Z).$$

1007 Using Theorem 2.2 again, the conditional entropy can be expressed as

$$1008 \quad h(f(X) + Z \mid Z) = h(\mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) - D_{\text{KL}}(f(X) + Z \parallel \mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) - I(f(X) + Z; Z).$$

1009 To conclude it remains to calculate $h(\mathbf{U}([- \varepsilon; \varepsilon + 1]^d)) = d \log(1 + 2\varepsilon)$. □

1010 **Corollary B.6.** *In the uniform distribution matching setup (Theorem 3.5), we have*

$$1011 \quad D_{\text{KL}}(f(X) \parallel \mathbf{U}([0; 1]^d)) \leq I(Z; f(X) + Z) - I(f(X'); f(X) + Z) - d \log(2\varepsilon).$$

1012 *Proof.* Applying Lemma B.4 alongside Lemma B.5 is enough. □

Table 2: The NN architectures used to conduct the tests on MNIST images in Section 5.

NN	Architecture
ConvNet, 24 × 24 images	×1: Conv2d(1, 32, ks=3), MaxPool2d(2), BatchNorm2d, LeakyReLU(0.01)
	×1: Conv2d(32, 64, ks=3), MaxPool2d(2), BatchNorm2d, LeakyReLU(0.01)
	×1: Conv2d(64, 128, ks=3), MaxPool2d(2), BatchNorm2d, LeakyReLU(0.01)
	×1: Dense(128, 128), LeakyReLU(0.01), Dense(128, dim)
Critic NN, pairs of vectors	×1: Dense(dim + dim, 256), LeakyReLU(0.01)
	×1: Dense(256, 256), LeakyReLU(0.01), Dense(256, 1)

C DETAILS OF IMPLEMENTATION

For experiments on MNIST dataset, we use a simple ConvNet with three convolutional and two fully connected layers. A three-layer fully-connected perceptron serves as a critic network for the InfoNCE loss. We provide the details in Table 2. We use additive Gaussian noise with $\sigma = 0.6$ as an input augmentation. Training hyperparameters are as follows: batch size = 1024, 2000 epochs, Adam optimizer (Kingma & Ba, 2017) with learning rate 10^{-3} .

The results on CIFAR datasets (Krizhevsky, 2009) in Table 1 were obtained with the standard configuration of SSL methods. Namely, we use ResNet-18 (He et al., 2015) backbone. Projection head for SimCLR consists of two linear layers, for VICReg – 3 layers. Respective configurations are [2048, 256] and [2048, 2048, 2048], meaning embedding dimensions are 256 and 2048, respectively. We apply a standard set of augmentations:

```

PretrainTransform: Compose(
  RandomResizedCrop(
    size=(32, 32),
    scale=(0.08, 1.0),
    ratio=(0.75, 1.3333333333333333),
    interpolation=InterpolationMode.BICUBIC,
    antialias=True
  )
  RandomApply(
    ColorJitter(
      brightness=(0.6, 1.4),
      contrast=(0.6, 1.4),
      saturation=(0.8, 1.2),
      hue=(-0.1, 0.1)
    )
  )
  RandomGrayscale(p=0.2)
  GaussianBlur(p=0.0)
  Solarization(p=0.0)
  RandomHorizontalFlip(p=0.5)
  ToTensor()
  Normalize(
    mean=[0.4914, 0.4822, 0.4465],
    std=[0.247, 0.2435, 0.2616],
    inplace=False
  )
)
TestTransform: Compose(
  ToTensor()
  Normalize(
    mean=[0.4914, 0.4822, 0.4465],
    std=[0.247, 0.2435, 0.2616],

```

```

1080         inplace=False
1081     )
1082 )
1083

```

1084 Training hyperparameters are as follows: batch size 256, 800 epochs, LARS optimizer (You et al.,
 1085 2017) with clipping, base learning rate 0.3, momentum 0.9, trust coefficient 0.02, weight decay
 1086 10^{-4} . For SimCLR, we use temperature 0.2, for VICReg – standard hyperparameters (25, 25, 1).

1087 We also provide the source code in the supplementary materials.

1088

1089

1090 D CONDITIONING GENERATIVE ADVERSARIAL NETWORKS

1091

1092 In this section, we leverage our method to generate conditioning vectors for a conventional cGAN
 1093 setup (Mirza & Osindero, 2014). We use the two-dimensional Gaussian embeddings of the MNIST
 1094 dataset, acquired from the noise level $\sigma = 0.1$. For conditioned generation, we get embeddings from
 1095 a batch of original images. For unconditioned generation, embeddings are sampled from $\mathcal{N}(0, I)$.
 1096 The results are presented in Figures 3 and 4.

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

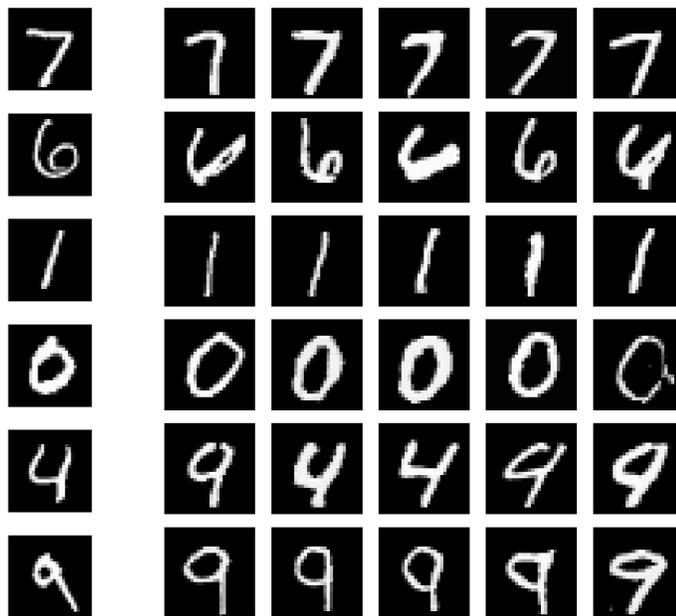
1129

1130

1131

1132

1133



(a) Encoded images

(b) Generated images

Figure 3: Results of conditional generation.

1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145
 1146
 1147
 1148
 1149
 1150
 1151
 1152
 1153
 1154
 1155
 1156
 1157
 1158
 1159
 1160
 1161
 1162
 1163
 1164
 1165
 1166
 1167
 1168
 1169
 1170
 1171
 1172
 1173
 1174
 1175
 1176
 1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

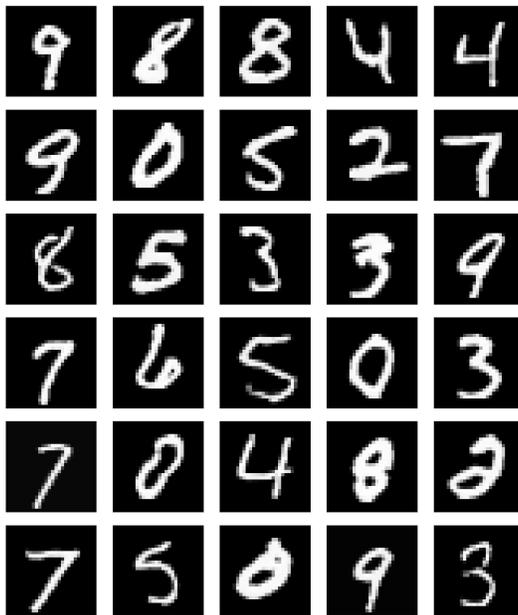


Figure 4: Results of unconditional generation.

E ADDITIONAL EXPERIMENTS ON CIFAR-10

Similarly to embedding MNIST into two-dimensional space, we run experiments with CIFAR-10. We use a ResNet-18 architecture augmented with an additional linear layer to map the 512-dimensional output to \mathbb{R}^2 , along with a batch normalization layer without learnable parameters (i.e., 'affine=False'). We experimented with both contrastive (SimCLR) and non-contrastive (VICReg) objectives, observing similar results across both methods. The experimental pipeline closely follows the theoretical setup:

- Unaugmented Input: The unaugmented image X is fed to the encoder, yielding $f(X)$, followed by noise injection to obtain $f(X) + Z$.
- Augmented Input: The augmented image X' is fed to the encoder to obtain $f(X')$, without noise injection.

Both outputs are then processed by the projection head and loss function (e.g., InfoNCE). The noise Z is sampled from a normal distribution with standard deviation σ . We conducted pre-training across various noise magnitudes with $\sigma \in \{0.0, 0.01, 0.025, 0.05, 0.1, 0.2, 0.2658, 0.5\}$. For each σ , we ran training with 5 different random seeds with `base_lr` = 0.1 for 400 epochs (visualized in Figure 5). To speed up convergence we additionally trained 3 models ($\sigma \in \{0, 0.05, 0.1\}$) with `base_lr` = 0.5 for 800 epochs (visualized in Figure 6).

After pre-training, we evaluated downstream performance using ROC AUC and performed normality tests on the learned 2D representations from a total of 40 trained models (that were trained for 400 epochs) (see Figure 7), depicting similar results to our previous MNIST experiments (Figure 1).

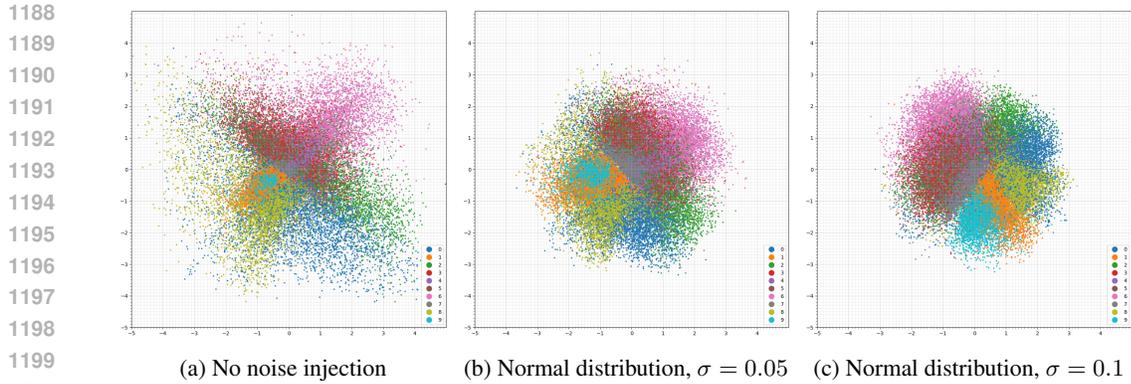


Figure 5: Visualization of 2D representations of the CIFAR-10 dataset, lr=0.1, 400 epochs.

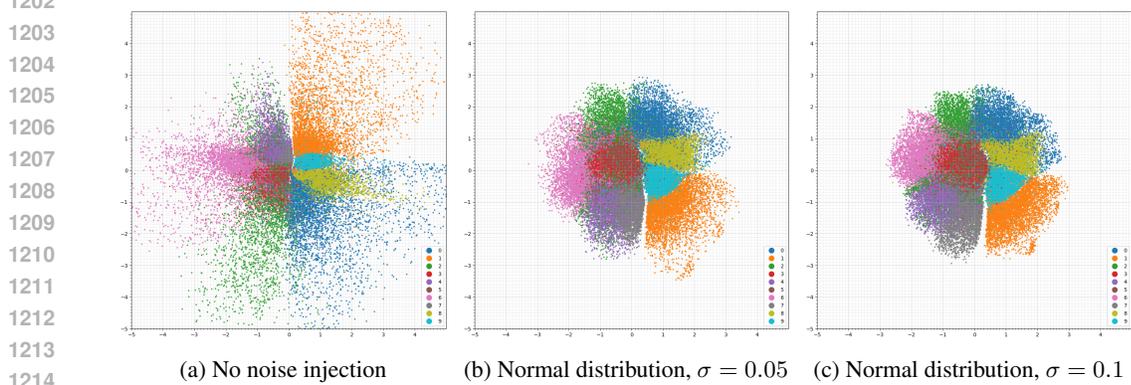


Figure 6: Visualization of 2D representations of the CIFAR-10 dataset, lr=0.5, 800 epochs.

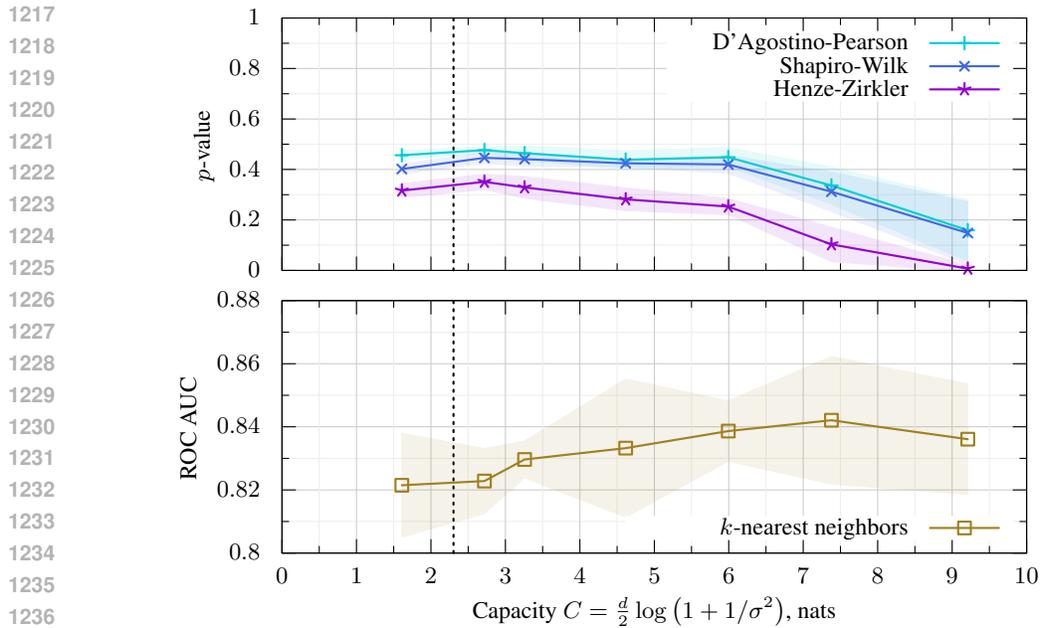


Figure 7: Results for CIFAR10 dataset in the Gaussian DM setup for $d = 2$ with varying capacity $C = \frac{d}{2} \log(1 + 1/\sigma^2)$, measured in nats (units of information based on natural logarithms). The dotted line denotes the minimal capacity required to preserve the information about the class labels in $f(X) + Z$. We run 5 experiments for each point and report mean values and 99% asymptotic confidence intervals. InfoNCE loss is used to approximate (2).

F ADDITIONAL EXPERIMENTS ON IMAGENET

To better assess the accuracy-DM trade-off in complex setups, additional experiments with the ImageNet datasets are conducted with VICReg loss.

For ImageNet-100, as backbone we use ResNet-18 followed by batch normalization layer without learnable parameters, i.e. `BatchNorm1d(512, affine=False)`. Projection head consists of standard MLP sequence: `nn.Linear(512, 2048), nn.BatchNorm1d(2048), nn.ReLU(), nn.Linear(2048, 2048), nn.BatchNorm1d(2048), nn.ReLU(), nn.Linear(2048, 2048)`.

For ImageNet-1k, we use ResNet-50 as backbone architecture, followed by batch normalization layer without learnable parameters (`dim=2048`). The projection head is comprised similar to the above setup of 3-layer MLP with hidden dimension = 8192.

In both setups, we run pre-training for 100 epochs and conduct linear probing afterwards closely following standardized protocol for augmentation and training. for ImageNet-1k, batch size and other (e.g. optimizer, data augmentation) hyperparameters set to specified in the original paper Bardes et al. (2021), however, evaluation batch size is set to 2048 for faster results. For ImageNet-100, batch size = 256, optimizer’s trust coefficient = 0.02, base learning rate = 0.5, weight decay = $1e-4$, inspired by da Costa et al. (2022). The resulting accuracies are reported in Table 3 across 5 random seeds for ImageNet-100 and on 1 seed on ImageNet-1k due to compute limitations.

	ImageNet-100		ImageNet	
	top-1	top-5	top-1	top-5
SimCLR	72.18 ± 0.40	92.02 ± 0.12	67.57	87.54
SimCLR $\sigma = 0.05$	72.27 ± 0.38	91.99 ± 0.18	67.33	87.54
SimCLR $\sigma = 0.1$	72.07 ± 0.27	91.65 ± 0.13	67.21	87.42
SimCLR $\sigma = 0.2$	71.68 ± 0.50	91.61 ± 0.24	67.17	87.30

Table 3: Noise magnitude influence on linear probing on ImageNet-100. Top-1 and top-5 accuracy (in %) is reported.