
High Mutual Information in Representation Learning with Symmetric Variational Inference

Micha Livne
University of Toronto
Vector Institute
mlivne@cs.toronto.edu

Kevin Swersky
Google Research
kswersky@google.com

David J. Fleet
University of Toronto
Vector Institute
fleet@cs.toronto.edu

Abstract

We introduce the Mutual Information Machine (MIM), a novel formulation of representation learning, using a joint distribution over the observations and latent state in an encoder/decoder framework. Our key principles are symmetry and mutual information, where symmetry encourages the encoder and decoder to learn different factorizations of the same underlying distribution, and mutual information, to encourage the learning of useful representations for downstream tasks. Our starting point is the symmetric Jensen-Shannon divergence between the encoding and decoding joint distributions, plus a mutual information encouraging regularizer. We show that this can be bounded by a tractable cross entropy loss function between the true model and a parameterized approximation, and relate this to the maximum likelihood framework. We also relate MIM to variational autoencoders (VAEs) and demonstrate that MIM is capable of learning symmetric factorizations, with high mutual information that avoids posterior collapse.

1 Introduction

The variational auto-encoder (VAE) [11] is a class of hierarchical Bayesian models based on transforming data from a simple latent prior into a complicated observation distribution using a neural network. The model density is represented by the marginal distribution, $\int_{\mathbf{z}} p_{\theta}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}) d\mathbf{z}$, where θ denotes the model parameters. The data log-likelihood is in general intractable to compute, so VAEs maximize the so-called evidence lower bound (ELBO).

$$\log \mathcal{P}(\mathbf{x}) \geq \mathbb{E}_{\mathbf{z} \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\log p_{\theta}(\mathbf{x}|\mathbf{z})] - \mathcal{D}_{\text{KL}}(q_{\theta}(\mathbf{z}|\mathbf{x}) \parallel \mathcal{P}(\mathbf{z})) , \quad (1)$$

where $q_{\theta}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})$ denotes an approximate posterior that maps observations to a distribution over latent variables, and where $p_{\theta}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})$ denotes the mapping from latent codes to observations. While VAEs have been successfully applied in a number of settings [8, 14, 12, 22], they suffer from a shortcoming known as posterior collapse [4, 5, 19, 24, 25]. Effectively, the penalty on the prior term in Eq. (1) causes the encoder to lose information in some dimensions of the latent code, leading the decoder to ignore those dimensions. The result is a model that generates good samples, but suffers from a poor latent representation. This is an issue for any downstream application that relies on \mathbf{z} .

Here we present a related but distinct framework called Mutual Information Machine (MIM). Our objectives are twofold: to promote symmetry in the encoding and decoding distributions (*i.e.*, consistency), and to encourage high mutual information between \mathbf{x} and \mathbf{z} (*i.e.*, good representation). By symmetry, we mean that the encoding and decoding distributions should represent two equivalent factorizations of the same underlying joint distribution. The framework can be seen as a symmetric analogue of VAEs, where we optimize a symmetric divergence instead of the asymmetric KL. In preliminary experiments, we find that MIM produces highly informative representations, with comparable sample quality to VAEs.

2 Mutual Information Machine

To begin, it is helpful to write down the VAE loss function as the KL divergence between two joint distributions over \mathbf{x} and \mathbf{z} , as in [18] i.e., up to an additive constant the following holds.

$$\mathcal{L}_{\text{VAE}}(\boldsymbol{\theta}) = \mathcal{D}_{\text{KL}}(q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})\mathcal{P}(\mathbf{x}) \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})\mathcal{P}(\mathbf{z})) . \quad (2)$$

This loss could be symmetrized by using the symmetric KL divergence, however this would result in a $\log \mathcal{P}(\mathbf{x})$ term that cannot be evaluated, as this distribution does not typically have an analytical form. In Pu et al. [18], minimization of the symmetric KL loss is expressed as a stationary point of an adversarial min-max objective.

Here we propose to optimize a bound on a regularized version of the Jensen-Shannon divergence (JSD). Defining the mixture distribution $\mathcal{M}_S \equiv \frac{1}{2}(p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})\mathcal{P}(\mathbf{z}) + q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})\mathcal{P}(\mathbf{x}))$, the JSD is defined as

$$\text{JSD}(\boldsymbol{\theta}) = \frac{1}{2} \left(\mathcal{D}_{\text{KL}}(p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})\mathcal{P}(\mathbf{z}) \| \mathcal{M}_S) + \mathcal{D}_{\text{KL}}(q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})\mathcal{P}(\mathbf{x}) \| \mathcal{M}_S) \right) . \quad (3)$$

A second principle of our formulation involves encouraging a representation with high mutual information. Mutual information is related to joint entropy by the identity $H(\mathbf{x}, \mathbf{z}) = H(\mathbf{x}) + H(\mathbf{z}) - I(\mathbf{x}; \mathbf{z})$. Since $H(\mathbf{x}, \mathbf{z})$ is tractable up to a constant under both the encoding and decoding distributions, we add the regularizer¹ $R_H(\boldsymbol{\theta}) = \frac{1}{2}(H_{q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})\mathcal{P}(\mathbf{x})}(\mathbf{x}, \mathbf{z}) + H_{p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})\mathcal{P}(\mathbf{z})}(\mathbf{x}, \mathbf{z}))$ to $\text{JSD}(\boldsymbol{\theta})$, which can be shown to be equal to $H_{\mathcal{M}_S}(\mathbf{x}, \mathbf{z})$, the entropy of the mixture distribution \mathcal{M}_S .

This formulation allows us to derive a tractable (in terms of being amenable to optimization via reparameterization) bound on the objective. First we define parameterized joint distributions, $q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) \equiv q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x})q_{\boldsymbol{\theta}}(\mathbf{x})$, and $p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) \equiv p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})p_{\boldsymbol{\theta}}(\mathbf{z})$. Where $p_{\boldsymbol{\theta}}(\mathbf{z})$ can be parameterized for added flexibility, or set to $\mathcal{P}(\mathbf{z})$. $q_{\boldsymbol{\theta}}(\mathbf{x})$ is an observation model which could be e.g., a normalizing flow [6, 13] or an auto-regressive model [23]. For a fair comparison with VAEs, we leverage the encoder and decoder distributions $q_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})$, $p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z})$ to avoid adding additional parameters. Details are given in the supplementary material.

Defining $\mathcal{M}_{\boldsymbol{\theta}} \equiv \frac{1}{2}(q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) + p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}))$, the bound can be expressed as,

$$H_{\mathcal{M}_S}(\mathbf{x}, \mathbf{z}) \leq H(\mathcal{M}_S, \mathcal{M}_{\boldsymbol{\theta}}) \leq \frac{1}{2} (H(\mathcal{M}_S, q_{\boldsymbol{\theta}}) + H(\mathcal{M}_S, p_{\boldsymbol{\theta}})) \equiv \mathcal{L}_{\text{MIM}}(\boldsymbol{\theta}) \quad (4)$$

where $H(p, q)$ is the cross entropy between distributions p and q . While the first inequality no longer depends on $\log \mathcal{P}(\mathbf{x})$ and is therefore tractable, the second inequality, which follows from Jensen's inequality, can be shown to additionally encourage consistency in $\mathcal{M}_{\boldsymbol{\theta}}$ (i.e., $q_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}) = p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})$). Further, when $q_{\boldsymbol{\theta}}(\mathbf{x}) = \mathcal{P}(\mathbf{x})$ and $p_{\boldsymbol{\theta}}(\mathbf{z}) = \mathcal{P}(\mathbf{z})$, it can be shown that $\mathcal{L}_{\text{MIM}}(\boldsymbol{\theta})$ is equal to the symmetric KL loss plus the regularizer $R_H(\boldsymbol{\theta})$. See supplementary material for more details.

We have therefore derived a loss function that encourages both symmetry in the encoding and decoding distributions, as well as high mutual information in the learned representation. Further, this loss function can be directly minimized without requiring an adversarial reformulation. In the experiments, we will show preliminary results exploring the properties of MIM.

3 Experiment: 2D Mixture Model Data

We begin with a dataset of 2D observations $\mathbf{x} \in \mathbb{R}^2$ drawn from a Gaussian mixture model, and a 2D latent space, $\mathbf{z} \in \mathbb{R}^2$. In 2D we can easily visualize the model and measure quantitative properties of interest (e.g., mutual information). (Complete experiment details given in supplementary material).

Figure 1 depicts results for the VAE (even columns) and MIM (odd columns) using single-layer encoder and decoder networks, with increasing numbers of hidden units (moving left to right) to control model expressiveness. The top row (for VAE and MIM respectively) depicts observation space. With each case we also report the mutual information and the root-mean-squared observation reconstruction error when sampling the predictive encoder/decoder distributions, with MIM showing a superior performance. See additional results in the supplementary material (Fig. 3).

The bottom row of Fig. 1 depicts the latent space behavior. For the weakest architecture, with only 5 hidden units, both MIM and VAE posteriors have large variances. When the number of hidden units increases, however, it is clear that while the VAE posterior variance remains very large in one

¹We use $H_q(\mathbf{x}, \mathbf{z})$ and $MI_q(\mathbf{x}; \mathbf{z})$ to notate entropy and mutual information under a distribution $q(\mathbf{x}, \mathbf{z})$

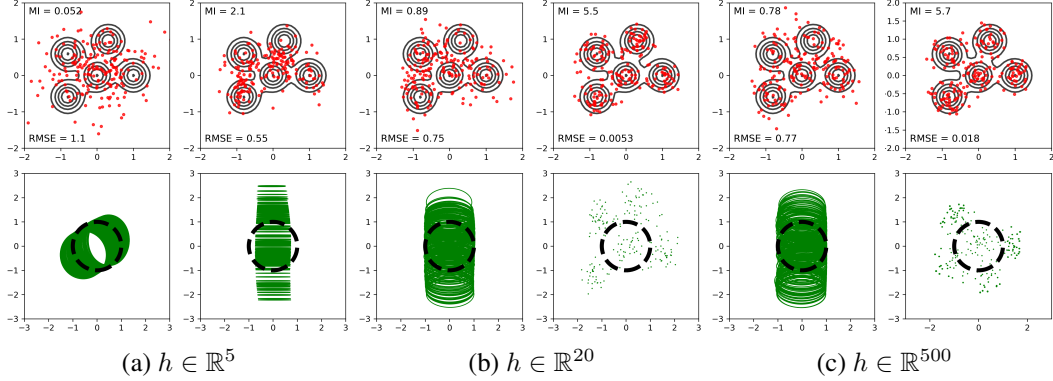


Figure 1: VAE (odd columns) and MIM (even columns) with 2D inputs, a 2D latent space, and 5, 20 and 500 hidden units. Top: Black contours are level sets of data distribution $\mathcal{P}(\mathbf{x})$, red points are reconstructed samples drawn from the decoder $p_{\theta}(\mathbf{x}|\mathbf{z}')$, where $\mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x}')$ for data point \mathbf{x}' from $\mathcal{P}(\mathbf{x})$. Bottom: The dashed black circle depicts one standard deviation of $\mathcal{P}(\mathbf{z})$. Each green curve depicts a one standard deviation ellipse of the encoder posterior $q_{\theta}(\mathbf{z}|\mathbf{x}')$. (a) For weak architectures MIM and VAE exhibit high posterior variance. (b,c) For more expressive architectures the VAE predictive variance remains high, an indication of posterior collapse. MIM generally produces lower predictive variance and lower reconstruction errors, consistent with high mutual information.

dimension (*i.e.*, a common sign of posterior collapse), the MIM encoder produces much tighter posteriors densities, which capture the global (*i.e.*, aggregated) structure of the observations.

In addition, with a more expressive architecture, *i.e.*, more hidden units, the MIM encoding variance is extremely small, and the reconstruction error approaches 0. In effect, the encoder and decoder learn an (approximately) invertible mapping using an unconstrained architecture (demonstrated here in 2D), when the dimensionality of the latent representation and the observations is the same.

The VAE, by comparison, is prone to posterior collapse, reflected in relatively low mutual information. In this regard, we note that several papers have described ways to mitigate posterior collapse in VAE learning, *e.g.*, by lower bounding or annealing the KL divergence term in the VAE objective (*e.g.*, [1, 19]), or by limiting the expressiveness of the decoder (*e.g.*, [5]). We posit that MIM does not suffer from this problem as a consequence of the objective design principles that encourage high mutual information between observations and the latent representation.

4 Experiment: MIM Representations with High Dimensional Image Data




Figure 2: VAE (odd) and MIM (even) \mathbf{z} embedding for Fashion MNIST and MNIST with convHVAE (S) architecture (see Table 1).

Dataset	convHVAE (S)		convHVAE (VP)	
	MIM	VAE	MIM	VAE
Fashion MNIST	0.84	0.75	0.84	0.76
MNIST	0.98	0.91	0.98	0.85

Table 1: Test classification accuracy of 5-NN classifier for MIM and VAE learning. MIM shows superior clustering of classes in the latent representation in an unsupervised manner.

Here we explore learning on higher dimensional image data, where we cannot accurately estimate mutual information [2]. Instead, following [9], we focus on an auxiliary classification task as a proxy for the quality of the learned representation and on qualitative visualization of it. We experiment with MNIST [16], and Fashion MNIST [27]. In what follows we also explore multiple architectures of VAE models from [21], and the corresponding MIM models (see Algorithm. 1 in the supplementary material), where, again we use VAE as the baseline. See supplementary material for the details.

For the auxiliary transfer learning classification task we opted for K-NN classification, being a non-parametric method which represents the clustering in the latent representation without any additional training. We show quantitative results in Table 1 for K-NN classification ($k = 5$). We also present the corresponding qualitative visual clustering results (*i.e.*, projection to 2D using t-SNE [26]) in Fig. 2. Here, it is clear that MIM learning tends to cluster classes in the latent representation better than VAE, for an identical parameterization of a model.

References

- [1] Alexander A. Alemi, Ben Poole, Ian Fischer, Joshua V. Dillon, Rif A. Saurous, and Kevin Murphy. An information-theoretic analysis of deep latent-variable models. *CoRR*, abs/1711.00464, 2017.
- [2] Ishmael Belghazi, Sai Rajeswar, Aristide Baratin, R Devon Hjelm, and Aaron Courville. MINE: Mutual information neural estimation. In *ICML*, 2018.
- [3] Jörg Bornschein, Samira Shabanian, Asja Fischer, and Yoshua Bengio. Bidirectional helmholtz machines. *CoRR*, abs/1506.03877, 2015.
- [4] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Józefowicz, and Samy Bengio. Generating sentences from a continuous space. *CoRR*, abs/1511.06349, 2015.
- [5] Xi Chen, Diederik P. Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. *CoRR*, abs/1611.02731, 2016.
- [6] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *International Conference on Learning Representations*, 2017.
- [7] Weihao Gao, Sewoong Oh, and Pramod Viswanath. Demystifying fixed k-nearest neighbor information estimators. *CoRR*, abs/1604.03006, 2016. URL <http://arxiv.org/abs/1604.03006>.
- [8] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: Masked autoencoder for distribution estimation. In *ICML*, 2015.
- [9] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019.
- [10] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv e-prints*, art. arXiv:1412.6980, Dec 2014.
- [11] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *ICLR*, 2013.
- [12] Diederik P. Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improving variational inference with inverse autoregressive flow. In *NIPS*, 2016.
- [13] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- [14] Jack Klys, Jake Snell, and Richard Zemel. Learning latent subspaces in variational autoencoders. In *NIPS*, 2018.
- [15] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Phys. Rev. E*, 69:066138, Jun 2004. doi: 10.1103/PhysRevE.69.066138. URL <https://link.aps.org/doi/10.1103/PhysRevE.69.066138>.
- [16] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2324, 1998.
- [17] Aaron van den Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, and George van den Driessche. Parallel wavenet: Fast high-fidelity speech synthesis. *arXiv preprint arXiv:1711.10433*, 2017.
- [18] Yuchen Pu, Weiyao Wang, Ricardo Henao, Liqun Chen, Zhe Gan, Chunyuan Li, and Lawrence Carin. Adversarial symmetric variational autoencoder. In *Advances in Neural Information Processing Systems*, pages 4330–4339, 2017.
- [19] Ali Razavi, Aäron van den Oord, Ben Poole, and Oriol Vinyals. Preventing posterior collapse with delta-vaes. *CoRR*, abs/1901.03416, 2019.
- [20] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative Models. In *ICML*, 2014.

- [21] Jakub M. Tomczak and Max Welling. VAE with a vampprior. *CoRR*, abs/1705.07120, 2017. URL <http://arxiv.org/abs/1705.07120>.
- [22] Rianne van den Berg, Leonard Hasenclever, Jakub M. Tomczak, and Max Welling. Sylvester normalizing flows for variational inference. *arXiv:1803.05649*, 2018.
- [23] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016.
- [24] Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *CoRR*, abs/1601.06759, 2016.
- [25] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. *CoRR*, abs/1711.00937, 2017.
- [26] Laurens van der Maaten and Geoffrey E. Hinton. Visualizing high-dimensional data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [27] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *CoRR*, abs/1708.07747, 2017. URL <http://arxiv.org/abs/1708.07747>.

A Detail Derivation of MIM Learning

Here we provide a detailed derivation of the loss of MIM learning, as defined in (4). We would like to formulate a loss function which includes (3) that reflects our desire for model symmetry and high mutual information. This objective is difficult to optimize directly since we do not know how to evaluate $\log \mathcal{P}(\mathbf{x})$ in the general case (*i.e.*, we do not have an exact closed-form expression for $\mathcal{P}(\mathbf{x})$). As a consequence, we introduce parameterized approximate priors, $q_\theta(\mathbf{x})$ and $p_\theta(\mathbf{z})$, to derive tractable bounds on the penalized Jensen-Shannon divergence. This is similar in spirit to VAEs, which introduce a parameterized approximate posterior. These parameterized priors, together with the conditional encoder and decoder, $q_\theta(\mathbf{z}|\mathbf{x})$ and $p_\theta(\mathbf{x}|\mathbf{z})$, comprise a new pair of joint distributions,

$$\begin{aligned} q_\theta(\mathbf{x}, \mathbf{z}) &\equiv q_\theta(\mathbf{z}|\mathbf{x}) q_\theta(\mathbf{x}) \\ p_\theta(\mathbf{x}, \mathbf{z}) &\equiv p_\theta(\mathbf{x}|\mathbf{z}) p_\theta(\mathbf{z}) . \end{aligned}$$

These new joint distributions allow us to formulate a new, tractable loss that bounds $H(\mathcal{M}_S)$:

$$\begin{aligned} \mathcal{L}_{\text{CE}}(\theta) &\equiv H(\mathcal{M}_S, \mathcal{M}_\theta) \\ &= \mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel \mathcal{M}_\theta) + H(\mathcal{M}_S) \\ &\geq H(\mathcal{M}_S) , \end{aligned} \tag{5}$$

where $H(\mathcal{M}_S, \mathcal{M}_\theta)$ denotes the cross-entropy between \mathcal{M}_S and \mathcal{M}_θ , and

$$\mathcal{M}_\theta = \frac{1}{2} (p_\theta(\mathbf{x}, \mathbf{z}) + q_\theta(\mathbf{x}, \mathbf{z})) . \tag{6}$$

In what follows we refer to \mathcal{L}_{CE} as the cross-entropy loss. It aims to match the model prior distributions to the anchors, while also minimizing $H(\mathcal{M}_S)$. The main advantage of this formulation is that the cross-entropy loss can be trained by Monte Carlo sampling from the anchor distributions with the reparameterization trick [11, 20].

At this stage it might seem odd to introduce a parametric prior for $\mathcal{P}(\mathbf{z})$. Indeed, setting it directly is certainly an option. Nevertheless, in order to achieve consistency between $p_\theta(\mathbf{x}, \mathbf{z})$ and $q_\theta(\mathbf{x}, \mathbf{z})$ it can be advantageous to allow $p_\theta(\mathbf{z})$ to vary. Essentially, we trade-off latent prior fidelity for increased model consistency.

One issue with \mathcal{L}_{CE} is that, while it will try to enforce consistency between the model and the anchored distributions, *i.e.*, $p_\theta(\mathbf{x}, \mathbf{z}) \approx p_\theta(\mathbf{x}|\mathbf{z})\mathcal{P}(\mathbf{z})$ and $q_\theta(\mathbf{x}, \mathbf{z}) \approx q_\theta(\mathbf{z}|\mathbf{x})\mathcal{P}(\mathbf{x})$, it will not directly try to achieve model consistency: $p_\theta(\mathbf{x}, \mathbf{z}) \approx q_\theta(\mathbf{x}, \mathbf{z})$. To remedy this, we bound \mathcal{L}_{CE} using Jensen's inequality, *i.e.*,

$$\mathcal{L}_{\text{MIM}}(\theta) \equiv \frac{1}{2} (H(\mathcal{M}_S, q_\theta(\mathbf{x}, \mathbf{z})) + H(\mathcal{M}_S, p_\theta(\mathbf{x}, \mathbf{z}))) \tag{7}$$

$$\geq \mathcal{L}_{\text{CE}}(\theta) . \tag{8}$$

Equation (7) gives us the loss function for the Mutual Information Machine (MIM). It is an average of cross entropy terms between the mixture distribution \mathcal{M}_S and the model encoding and decoding distributions respectively. To see that this encourages model consistency, it can be shown that \mathcal{L}_{MIM} is equivalent to \mathcal{L}_{CE} plus a non-negative model consistency regularizer; *i.e.*,

$$\mathcal{L}_{\text{MIM}}(\theta) = \mathcal{L}_{\text{CE}}(\theta) + \text{R}_{\text{MIM}}(\theta) . \tag{9}$$

The non-negativity of R_{MIM} is a simple consequence of $\mathcal{L}_{\text{MIM}}(\theta) \geq \mathcal{L}_{\text{CE}}(\theta)$ in (8).

In what follows we derive the form of the MIM consistency regularizer in Eq. (4), named $\text{R}_{\text{MIM}}(\theta)$. Recall that we define $\mathcal{M}_\theta = \frac{1}{2}(p_\theta(\mathbf{x}, \mathbf{z}) + q_\theta(\mathbf{x}, \mathbf{z}))$. We can show that \mathcal{L}_{MIM} is equivalent to $H(\mathcal{M}_S, \mathcal{M}_\theta)$ plus a regularizer by taking their difference.

$$\begin{aligned} \text{R}_{\text{MIM}}(\theta) &= \mathcal{L}_{\text{MIM}}(\theta) - H(\mathcal{M}_S, \mathcal{M}_\theta) \\ &= \frac{1}{2} (H(\mathcal{M}_S, p_\theta(\mathbf{x}, \mathbf{z})) + H(\mathcal{M}_S, q_\theta(\mathbf{x}, \mathbf{z}))) - H(\mathcal{M}_S, \mathcal{M}_\theta) \\ &= \frac{1}{2} (\mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel p_\theta(\mathbf{x}, \mathbf{z})) + H_{\mathcal{M}_S}(\mathbf{x}, \mathbf{z}) + \mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel q_\theta(\mathbf{x}, \mathbf{z})) + H_{\mathcal{M}_S}(\mathbf{x}, \mathbf{z})) \\ &\quad - \mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel \mathcal{M}_\theta) - H_{\mathcal{M}_S}(\mathbf{x}, \mathbf{z}) \\ &= \frac{1}{2} (\mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel p_\theta(\mathbf{x}, \mathbf{z})) + \mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel q_\theta(\mathbf{x}, \mathbf{z}))) - \mathcal{D}_{\text{KL}}(\mathcal{M}_S \parallel \mathcal{M}_\theta) \end{aligned} \tag{10}$$

where $R_{\text{MIM}}(\theta)$ is non-negative, and is zero only when the two joint model distributions, $q_\theta(\mathbf{x}, \mathbf{z})$ and $p_\theta(\mathbf{x}, \mathbf{z})$, are identical under fair samples from the joint sample distribution $\mathcal{M}_S(\mathbf{x}, \mathbf{z})$. To prove that $R_{\text{MIM}}(\theta) \geq 0$ we now construct Equation (10) in terms of expectation over a joint distribution, which yields

$$\begin{aligned} R_{\text{MIM}}(\theta) &= \frac{1}{2}(H(\mathcal{M}_S, p_\theta(\mathbf{x}, \mathbf{z})) + H(\mathcal{M}_S, q_\theta(\mathbf{x}, \mathbf{z}))) - H(\mathcal{M}_S, \mathcal{M}_\theta) \\ &= \mathbb{E}_{\mathbf{x}, \mathbf{z} \sim \mathcal{M}_S} \left[-\frac{1}{2} \log p_\theta(\mathbf{x}, \mathbf{z}) - \frac{1}{2} \log q_\theta(\mathbf{x}, \mathbf{z}) + \log \frac{1}{2}(q_\theta(\mathbf{x}, \mathbf{z}) + p_\theta(\mathbf{x}, \mathbf{z})) \right] \\ &= \mathbb{E}_{\mathbf{x}, \mathbf{z} \sim \mathcal{M}_S} \left[-\log \sqrt{q_\theta(\mathbf{x}, \mathbf{z}) \cdot p_\theta(\mathbf{x}, \mathbf{z})} + \log \frac{1}{2}(q_\theta(\mathbf{x}, \mathbf{z}) + p_\theta(\mathbf{x}, \mathbf{z})) \right] \\ &= \mathbb{E}_{\mathbf{x}, \mathbf{z} \sim \mathcal{M}_S} \left[-\log \frac{\sqrt{q_\theta(\mathbf{x}, \mathbf{z}) \cdot p_\theta(\mathbf{x}, \mathbf{z})}}{\frac{1}{2}(q_\theta(\mathbf{x}, \mathbf{z}) + p_\theta(\mathbf{x}, \mathbf{z}))} \right] \geq 0 \end{aligned}$$

where the inequality follows Jensen's inequality, and equality holds only when $q_\theta(\mathbf{x}, \mathbf{z}) = p_\theta(\mathbf{x}, \mathbf{z})$ (*i.e.*, encoding and decoding distributions are consistent). In practice we find that encouraging model consistency also helps stabilize learning.

To understand the MIM objective in greater depth, we find it helpful to express \mathcal{L}_{MIM} as a sum of fundamental terms that provide some intuition for its expected behavior. In particular, as derived in the supplementary material:

$$\begin{aligned} \mathcal{L}_{\text{MIM}}(\theta) &= R_H(\theta) + \frac{1}{4}(\mathcal{D}_{\text{KL}}(\mathcal{P}(z) \parallel p_\theta(z)) + \mathcal{D}_{\text{KL}}(\mathcal{P}(x) \parallel q_\theta(x))) \\ &\quad + \frac{1}{4}(\mathcal{D}_{\text{KL}}(q_\theta(z|x) \mathcal{P}(x) \parallel p_\theta(\mathbf{x}, \mathbf{z})) + \mathcal{D}_{\text{KL}}(p_\theta(\mathbf{x}|z) \mathcal{P}(z) \parallel q_\theta(\mathbf{z}, x))) \quad (11) \end{aligned}$$

The first term in (11), as discussed above, encourages high mutual information between observations and latent states. The second term shows that MIM directly encourages the model prior distributions to match the anchor distributions. Indeed, the KL term between the data anchor and the model prior is the maximum likelihood objective. The third term encourages consistency between the model distributions and the anchored distributions, in effect fitting the model decoder to samples drawn from the anchored encoder (cf. VAE), and, via symmetry, fitting the model encoder to samples drawn from the anchored decoder (both with reparameterization). In this view, MIM can be seen as simultaneously training and distilling a model distribution over the data into a latent variable model. The idea of distilling density models has been used in other domains, *e.g.*, for parallelizing auto-regressive models [17].

In summary, the MIM loss can be viewed as an upper bound on the entropy of a particular mixture distribution \mathcal{M}_S :

$$\begin{aligned} \mathcal{L}_{\text{MIM}}(\theta) &= \frac{1}{2}(H(\mathcal{M}_S, q_\theta(\mathbf{x}, \mathbf{z})) + H(\mathcal{M}_S, p_\theta(\mathbf{x}, \mathbf{z}))) \\ &= H(\mathcal{M}_S, \mathcal{M}_\theta) + R_{\text{MIM}}(\theta) \\ &\geq H(\mathcal{M}_S, \mathcal{M}_\theta) \\ &\geq H(\mathcal{M}_S) \\ &= H_{\mathcal{M}_S}(\mathbf{x}) + H_{\mathcal{M}_S}(\mathbf{z}) - I_{\mathcal{M}_S}(\mathbf{x}; \mathbf{z}) \quad (12) \end{aligned}$$

Through the MIM loss and the introduction of the parameterized model distribution \mathcal{M}_θ , we are pushing down on the entropy of the anchored mixture distribution \mathcal{M}_S , which is the sum of marginal entropies minus the mutual information. Minimizing the MIM bound yields consistency of the model encoder and decoder, and high mutual information of \mathcal{M}_S between observations and latent states.

B MIM in terms of Symmetric KL Divergence

As given in Equation (2), the VAE objective can be expressed as minimizing the KL divergence between the joint anchored encoding and anchored decoding distributions (*i.e.*, which jointly defines the sample distribution $\mathcal{M}_S(\mathbf{x}, \mathbf{z})$). Here we refer to $\mathcal{P}(\mathbf{x})$ and $\mathcal{P}(\mathbf{z})$ as anchors which are given

externally and are not learned. Below we consider a model formulation using the symmetric KL divergence (SKL),

$$\text{SKL}(\boldsymbol{\theta}) = \frac{1}{2} (\mathcal{D}_{\text{KL}} (p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}) \| q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}} (q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}) \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}))) ,$$

the second term of which is the VAE objective.

In what follows we explore the relation between SKL, JSD, and MIM. Recall that the JSD is written as,

$$\text{JSD}(\boldsymbol{\theta}) = \frac{1}{2} \left(\mathcal{D}_{\text{KL}} (p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}) \| \mathcal{M}_S) + \mathcal{D}_{\text{KL}} (q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}) \| \mathcal{M}_S) \right) .$$

Using the identity $\mathcal{D}_{\text{KL}} (p \| q) = H(p, q) - H_p(\mathbf{x}, \mathbf{z})$, we can express the JSD in terms of entropy and cross entropy.

$$\begin{aligned} \text{JSD}(\boldsymbol{\theta}) &= \frac{1}{2} \left(H(p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}), \mathcal{M}_S) - H_{p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})}(\mathbf{x}, \mathbf{z}) \right. \\ &\quad \left. + H(q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}), \mathcal{M}_S) - H_{q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})}(\mathbf{x}, \mathbf{z}) \right) \\ &= \frac{1}{2} (H(p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}), \mathcal{M}_S) + H(q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}), \mathcal{M}_S)) - \text{R}_H(\boldsymbol{\theta}) \end{aligned}$$

Using Jensen's inequality, we can bound $\text{JSD}(\boldsymbol{\theta})$ from above,

$$\begin{aligned} \text{JSD}(\boldsymbol{\theta}) &\leq \frac{1}{4} \left(H_{p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})}(\mathbf{x}, \mathbf{z}) + H(p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}), q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) \right. \\ &\quad \left. + H_{q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})}(\mathbf{x}, \mathbf{z}) + H(q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}), p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})) \right) - \text{R}_H(\boldsymbol{\theta}) \quad (13) \\ &= \frac{1}{4} \left(H(p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}), q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + H(q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}), p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})) \right. \\ &\quad \left. + 2\text{R}_H(\boldsymbol{\theta}) \right) - \text{R}_H(\boldsymbol{\theta}) \\ &= \frac{1}{4} \left(\mathcal{D}_{\text{KL}} (p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}) \| q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}} (q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}) \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})) \right. \\ &\quad \left. + 4\text{R}_H(\boldsymbol{\theta}) \right) - \text{R}_H(\boldsymbol{\theta}) \\ &= \frac{1}{4} (\mathcal{D}_{\text{KL}} (p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}) \| q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}} (q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}) \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}))) \\ &= \frac{1}{2} \text{SKL}(\boldsymbol{\theta}) \end{aligned}$$

If we add the regularizer $\text{R}_H(\boldsymbol{\theta})$ and combine terms, we get

$$\frac{1}{2} \text{SKL}(\boldsymbol{\theta}) + \text{R}_H(\boldsymbol{\theta}) = \frac{1}{2} (H(\mathcal{M}_S, q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + H(\mathcal{M}_S, p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})))$$

When the model priors $q_{\boldsymbol{\theta}}(\mathbf{x})$ and $p_{\boldsymbol{\theta}}(\mathbf{z})$ are equal to the fixed priors $\mathcal{P}(\mathbf{x})$ and $\mathcal{P}(\mathbf{z})$, this regularized SKL and MIM are equivalent. In general, however, the MIM loss is not a bound on the regularized SKL.

In what follows, we derive the exact relationship between JSD and SKL.

$$\begin{aligned} \frac{1}{2} \text{SKL}(\boldsymbol{\theta}) + \text{R}_H(\boldsymbol{\theta}) &= \frac{1}{2} (\mathcal{D}_{\text{KL}} (\mathcal{M}_S \| q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}} (\mathcal{M}_S \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}))) + H_{\mathcal{M}_S}(\mathbf{x}, \mathbf{z}) \\ &= \frac{1}{2} (\mathcal{D}_{\text{KL}} (\mathcal{M}_S \| q_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}} (\mathcal{M}_S \| p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}))) \\ &\quad + \text{JSD}(\boldsymbol{\theta}) + \text{R}_H(\boldsymbol{\theta}) \end{aligned}$$

which gives the exact relation between JSD and SKL.

$$\begin{aligned} \frac{1}{2}\text{SKL}(\theta) &= \frac{1}{2} (\mathcal{D}_{\text{KL}}(\mathcal{M}_S \| q_\theta(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}}(\mathcal{M}_S \| p_\theta(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}))) + \text{JSD}(\theta) \\ &= \frac{1}{2} (\mathcal{D}_{\text{KL}}(\mathcal{M}_S \| q_\theta(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x})) + \mathcal{D}_{\text{KL}}(\mathcal{M}_S \| p_\theta(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}))) \\ &\quad + \frac{1}{2} (\mathcal{D}_{\text{KL}}(q_\theta(\mathbf{z}|\mathbf{x}) \mathcal{P}(\mathbf{x}) \| \mathcal{M}_S) + \mathcal{D}_{\text{KL}}(p_\theta(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z}) \| \mathcal{M}_S)) \end{aligned}$$

C Parameterizing $q_\theta(\mathbf{x})$ and $p_\theta(\mathbf{z})$ for fair comparison with VAEs

In the MIM framework, there is flexibility in the choice of $q_\theta(\mathbf{x})$ and $p_\theta(\mathbf{z})$. To facilitate a direct comparison with VAEs, we must be careful to keep the architectures consistent and not introduce additional model parameters. For simplicity and a fair comparison, we set $p_\theta(\mathbf{z}) = \mathcal{P}(\mathbf{z})$ and leave other considerations for future work. For $q_\theta(\mathbf{x})$, we consider two different approaches that leverage the decoder distribution $p_\theta(\mathbf{x}|\mathbf{z})$. The first is to consider the marginal decoding distribution,

$$q_\theta(\mathbf{x}) = \mathbb{E}_{\mathcal{P}(\mathbf{z})} [p_\theta(\mathbf{x}|\mathbf{z})]. \quad (14)$$

Which we approximate by drawing one sample from $\mathcal{P}(\mathbf{z})$ when we need to evaluate $q_\theta(\mathbf{x})$. This can suffer from high variance if the prior is far from the true posterior.

The other is to consider an importance sampling estimate, for which we use the encoder distribution $q_\theta(\mathbf{x}|\mathbf{z})$,

$$q_\theta(\mathbf{x}) = \mathbb{E}_{q_\theta(\mathbf{z}|\mathbf{x})} \left[\frac{p_\theta(\mathbf{x}|\mathbf{z}) \mathcal{P}(\mathbf{z})}{q_\theta(\mathbf{z}|\mathbf{x})} \right]. \quad (15)$$

Where once again, we approximate using one sample, this time from the encoder distribution, and multiply by the importance weight $\frac{\mathcal{P}(\mathbf{z})}{q_\theta(\mathbf{z}|\mathbf{x})}$. Samples of \mathbf{z} are drawn using the reparameterization trick [11, 20] in order to allow for gradient-based training. We utilize (14) when sampling from the decoding distribution during the training of a MIM model, and (15) when sampling from the encoding distribution.

D Experimentation Details

Following [2], we estimate mutual information using the KSG mutual information estimator [15, 7], based on a K-NN neighborhoods with $k = 5$, and measure the quality of the representation with classification auxiliary task.

The learning algorithm is described in Algorithm 1. In what follows we describe in details the experimental setup, architecture, and training procedure for the experiments that were presented in the paper.

D.1 2D Mixture Model Data

In all experiments we use Adam optimizer [10] with $lr = 1e - 3$, and mini-batch of size 128. We stopped training for all experiments when validation loss has not improved for 10 epochs.

Data are drawn from a Gaussian mixture model with five isotropic components with standard deviation 0.25, and the latent anchor, $\mathcal{P}(\mathbf{z})$, is an isotropic standard normal distribution. The encoder and decoder are conditional Gaussian distributions, where the means and variances of which are regressed from the input using two fully connected layers and \tanh activation function. Following [3], the parameterized data prior, $q_\theta(\mathbf{x})$, is defined to be the marginal of the decoding distribution, or explicitly $q_\theta(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim p_\theta(\mathbf{z})} [p_\theta(\mathbf{x}|\mathbf{z})]$, where the only model parameters are those of the encoder and decoder, and the encoding distribution $q_\theta(\mathbf{x}, \mathbf{z})$ is defined to be consistent with the decoding distribution $p_\theta(\mathbf{x}, \mathbf{z})$. As such we can learn models with MIM and VAE objective that share the same architectures and parameterizations.

Algorithm 1 MIM learning with marginal $q_\theta(\mathbf{x})$

Require: Samples from anchors $\mathcal{P}(\mathbf{x}), \mathcal{P}(\mathbf{z})$

Require: Define $q_\theta(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim p_\theta(\mathbf{z})} [p_\theta(\mathbf{x}|\mathbf{z})]$

```
1: while not converged do
2:   # Sample encoding distribution
3:    $D_{\text{enc}} \leftarrow \{\mathbf{x}_i, \mathbf{z}_i \sim q_\theta(\mathbf{z}|\mathbf{x})\mathcal{P}(\mathbf{x})\}_{i=1}^N$ 
4:   # Compute objective, approximate  $\log q_\theta(\mathbf{x})$  with 1 sample
5:    $\log q_\theta(\mathbf{x}_i) \approx \log p_\theta(\mathbf{x}_i|\mathbf{z}_i) + \log p_\theta(\mathbf{z}_i) - \log q_\theta(\mathbf{z}_i|\mathbf{x}_i)$ 
6:    $\hat{\mathcal{L}}_{\text{MIM}}(\theta; D_{\text{enc}}) \leftarrow -\frac{1}{N} \sum_{i=1}^N (\log p_\theta(\mathbf{x}_i|\mathbf{z}_i) + \log p_\theta(\mathbf{z}_i))$ 
7:   # Sample decoding distribution
8:    $D_{\text{dec}} \leftarrow \{\mathbf{x}_i, \mathbf{z}_i \sim p_\theta(\mathbf{x}|\mathbf{z})\mathcal{P}(\mathbf{z})\}_{i=1}^N$ 
9:   # Compute objective, approximate  $\log q_\theta(\mathbf{x})$  with 1 sample and importance sampling
10:   $\log q_\theta(\mathbf{x}_i) \approx \log p_\theta(\mathbf{x}_i|\mathbf{z}_i)$ 
11:   $\hat{\mathcal{L}}_{\text{MIM}}(\theta; D_{\text{dec}}) \leftarrow -\frac{1}{2N} \sum_{i=1}^N (\log q_\theta(\mathbf{z}_i|\mathbf{x}_i) + 2 \log p_\theta(\mathbf{x}_i|\mathbf{z}_i) + \log p_\theta(\mathbf{z}_i))$ 
12:  # Minimize loss
13:   $\Delta\theta \propto -\nabla_\theta (\hat{\mathcal{L}}_{\text{MIM}}(\theta; D_{\text{dec}}) + \hat{\mathcal{L}}_{\text{MIM}}(\theta; D_{\text{enc}}))$ 
14: end while
```

D.2 Representation Learning with MIM in High Dimensional Image Data

We experiment with convHVAE ($L = 2$) model from [21], with Standard (S) prior which is the usual Normal distributions, and VampPrior (VP) prior which define the prior as a mixture model of the encoder conditioned on learnable pseudo-inputs \mathbf{u}_k , or explicitly $p_\theta(\mathbf{z}) = \frac{1}{K} \sum_{k=1}^K q_\theta(\mathbf{z}|\mathbf{u}_k)$. In all the experiments we used the same setup that was used in [21], and with the same latent dimensionality $\mathbf{z} \in \mathbb{R}^{80}$. By doing so we aim to highlight the generality of MIM learning as being architecture independent, and to provide examples for the training procedure of existing VAE architectures with MIM learning.

E Additional Results

E.1 2D Mixture Model Data




Figure 3: Test performance for MIM (blue) and VAE (red) for 2D GMM experiment, all as functions of the number of hidden units (on x-axis), based on 10 learned models in each case. From left to right, plots show mutual information, log marginal probability of test points, reconstruction error, and k-NN classification performance.

Here we quantify the complete experimental results that were presented in Fig. 1. We plot the mutual information, the average log marginal of test points under the model q_θ , the reconstruction error of test points, and 5-NN classification (predicting which of five GMM components the test points were drawn from).

E.2 Representation Learning with MIM in High Dimensional Image Data

Training times of MIM models are comparable to training times for VAEs with comparable architectures. The principal difference will be the time required for sampling from the decoder during




Figure 4: VAE and A-MIM learning with PixelHVAE (VP) for Fashion MNIST dataset. Top three rows are data samples, VAE, A-MIM, correspondingly. Bottom row is model samples.

training. For certain models, such as auto-regressive decoders [12], this can be significant. In such cases (*i.e.*, PixelHVAE here), we find that we can also learn the model by changing the sampling distribution to only include samples from the encoding distribution. By using asymmetric sampling $\mathcal{M}_S(\mathbf{x}, \mathbf{z}) = \mathcal{P}(\mathbf{x})q_\theta(\mathbf{z})$, where we sample from the encoding distribution only (*i.e.*, similar to VAE), training time is comparable to VAE. We name that model A-MIM.

Here we show qualitative results for the most expressive model, PixelHVAE (VP). Figures (4, 5), depict reconstruction, and sampling for Fashion-MNIST, and MNIST, correspondingly. The top three rows of each of the plots depicts data samples, VAE reconstruction, and A-MIM reconstruction, respectively. The bottom row depicts samples. The results demonstrate comparable samples and reconstruction for MIM and VAE.




Figure 5: VAE and A-MIM learning with PixelHVAE (VP) for MNIST dataset. Top three rows are data samples, VAE, A-MIM, correspondingly. Bottom row is model samples.