

Disentangling by Factorising

Hyunjik Kim^{1,2} Andriy Mnih¹

Abstract

We define and address the problem of unsupervised learning of disentangled representations on data generated from independent factors of variation. We propose FactorVAE, a method that disentangles by encouraging the distribution of representations to be factorial and hence independent across the dimensions. We show that it improves upon β -VAE by providing a better trade-off between disentanglement and reconstruction quality. Moreover, we highlight the problems of a commonly used disentanglement metric and introduce a new metric that does not suffer from them.

1. Introduction

Learning interpretable representations of data that expose semantic meaning has important consequences for artificial intelligence. Such representations are useful not only for standard downstream tasks such as supervised learning and reinforcement learning, but also for tasks such as transfer learning and zero-shot learning where humans excel but machines struggle (Lake et al., 2016). There have been multiple efforts in the deep learning community towards learning factors of variation in the data, commonly referred to as learning a *disentangled representation*. While there is no canonical definition for this term, we adopt the one due to Bengio et al. (2013): a representation where a change in one dimension corresponds to a change in one factor of variation, while being relatively invariant to changes in other factors. In particular, we assume that the data has been generated from a fixed number of independent factors of variation.³ We focus on image data, where the effect of factors of variation is easy to visualise.

Using generative models has shown great promise in learning disentangled representations in images. Notably, semi-

¹DeepMind, UK ²Department of Statistics, University of Oxford. Correspondence to: Hyunjik Kim <hyunjikk@google.com>.

³We discuss the limitations of this assumption in Section 4.

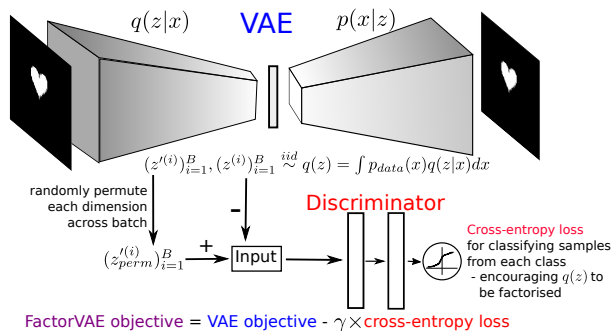


Figure 1. Architecture of FactorVAE, a Variational Autoencoder (VAE) that encourages the code distribution to be factorial. The top row is a VAE with convolutional encoder and decoder, and the bottom row is an MLP classifier, the discriminator, that distinguishes whether the input was drawn from the marginal code distribution or the product of its marginals.

supervised approaches that require implicit or explicit knowledge about the true underlying factors of the data have excelled at disentangling (Kulkarni et al., 2015; Kingma et al., 2014; Reed et al., 2014; Siddharth et al., 2017; Hinton et al., 2011; Mathieu et al., 2016; Goroshin et al., 2015; Hsu et al., 2017; Denton & Birodkar, 2017). However, ideally we would like to learn these in an unsupervised manner, due to the following reasons: 1. Humans are able to learn factors of variation unsupervised (Perry et al., 2010). 2. Labels are costly as obtaining them requires a human in the loop. 3. Labels assigned by humans might be inconsistent or leave out the factors that are difficult for humans to identify.

β -VAE (Higgins et al., 2016) is a popular method for unsupervised disentanglement based on the Variational Autoencoder (VAE) framework (Kingma & Welling, 2014; Rezende et al., 2014) for generative modelling. It uses a modified version of the VAE objective with a larger weight ($\beta > 1$) on the KL divergence between the variational posterior and the prior, and has proven to be an effective and stable method for disentanglement. One drawback of β -VAE is that reconstruction quality (compared to VAE) must be sacrificed in order to obtain better disentanglement. The goal of our work is to obtain a better trade-off between disentanglement and reconstruction, allowing to achieve better disentanglement without degrading reconstruction quality. In this work, we analyse the source of this trade-off and propose FactorVAE, which augments the VAE objective with a penalty that

encourages the marginal distribution of representations to be factorial without substantially affecting the quality of reconstructions. This penalty is expressed as a KL divergence between this marginal distribution and the product of its marginals, and is optimised using a discriminator network following the divergence minimisation view of GANs (Nowozin et al., 2016; Mohamed & Lakshminarayanan, 2016). Our experimental results show that this approach achieves better disentanglement than β -VAE for the same reconstruction quality. We also point out the weaknesses in the disentangling metric of Higgins et al. (2016), and propose a new metric that addresses these shortcomings.

A popular alternative to β -VAE is InfoGAN (Chen et al., 2016), which is based on the Generative Adversarial Net (GAN) framework (Goodfellow et al., 2014) for generative modelling. InfoGAN learns disentangled representations by rewarding the mutual information between the observations and a subset of latents. However at least in part due to its training stability issues (Higgins et al., 2016), there has been little empirical comparison between VAE-based methods and InfoGAN. Taking advantage of the recent developments in the GAN literature that help stabilise training, we include InfoWGAN-GP, a version of InfoGAN that uses Wasserstein distance (Arjovsky et al., 2017) and gradient penalty (Gulrajani et al., 2017), in our experimental evaluation.

In summary, we make the following contributions: 1) We introduce FactorVAE, a method for disentangling that gives higher disentanglement scores than β -VAE for the same reconstruction quality. 2) We identify the weaknesses of the disentanglement metric of Higgins et al. (2016) and propose a more robust alternative. 3) We give quantitative comparisons of FactorVAE and β -VAE against InfoGAN’s WGAN-GP counterpart for disentanglement.

2. Trade-off between Disentanglement and Reconstruction in β -VAE

We motivate our approach by analysing where the disentanglement and reconstruction trade-off arises in the β -VAE objective. First, we introduce notation and architecture of our VAE framework. We assume that observations $x^{(i)} \in \mathcal{X}$, $i = 1, \dots, N$ are generated by combining K underlying factors $f = (f_1, \dots, f_K)$. These observations are modelled using a real-valued latent/code vector $z \in \mathbb{R}^d$, interpreted as the representation of the data. The generative model is defined by the standard Gaussian prior $p(z) = \mathcal{N}(0, I)$, intentionally chosen to be a factorised distribution, and the decoder $p_\theta(x|z)$ parameterised by a neural net. The variational posterior for an observation is $q_\theta(z|x) = \prod_{j=1}^d \mathcal{N}(z_j|\mu_j(x), \sigma_j^2(x))$, with the mean and variance produced by the encoder, also parameterised by

a neural net.¹ The variational posterior can be seen as the distribution of the representation corresponding to the data point x . The distribution of representations for the entire data set is then given by

$$q(z) = \mathbb{E}_{p_{data}(x)}[q(z|x)] = \frac{1}{N} \sum_{i=1}^N q(z|x^{(i)}), \quad (1)$$

which is known as the marginal posterior or aggregate posterior, where p_{data} is the empirical data distribution. A disentangled representation would have each z_j correspond to precisely one underlying factor f_k . Since we assume that these factors vary independently, we wish for a factorial distribution $q(z) = \prod_{j=1}^d q(z_j)$.

The β -VAE objective

$$\frac{1}{N} \sum_{i=1}^N \left[\mathbb{E}_{q(z|x^{(i)})} [\log p(x^{(i)}|z)] - \beta KL(q(z|x^{(i)})||p(z)) \right]$$

is a variational lower bound on $\mathbb{E}_{p_{data}(x)}[\log p(x^{(i)})]$ for $\beta \geq 1$, reducing to the VAE objective for $\beta = 1$. Its first term can be interpreted as the negative *reconstruction error*, and the second term as the complexity penalty that acts as a regulariser. We may further break down this KL term as (Hoffman & Johnson, 2016; Makhzani & Frey, 2017)

$$\mathbb{E}_{p_{data}(x)}[KL(q(z|x)||p(z))] = I(x; z) + KL(q(z)||p(z)),$$

where $I(x; z)$ is the mutual information between x and z under the joint distribution $p_{data}(x)q(z|x)$. See Appendix C for the derivation. Penalising the $KL(q(z)||p(z))$ term pushes $q(z)$ towards the factorial prior $p(z)$, encouraging independence in the dimensions of z and thus disentangling. Penalising $I(x; z)$, on the other hand, reduces the amount of information about x stored in z , which can lead to poor reconstructions for high values of β (Makhzani & Frey, 2017). Thus making β larger than 1, penalising both terms more, leads to better disentanglement but reduces reconstruction quality. When this reduction is severe, there is insufficient information about the observation in the latents, making it impossible to recover the true factors. Therefore there exists a value of $\beta > 1$ that gives highest disentanglement, but results in a higher reconstruction error than a VAE.

3. Total Correlation Penalty and FactorVAE

Penalising $I(x; z)$ more than a VAE does might be neither necessary nor desirable for disentangling. For example, InfoGAN disentangles by encouraging $I(x; c)$ to be high where c is a subset of the latent variables z ². Hence we

¹In the rest of the paper we will omit the dependence of p and q on their parameters θ for notational convenience.

²Note however that $I(x; z)$ in β -VAE is defined under the joint distribution of data and their encoding distribution $p_{data}(x)q(z|x)$, whereas $I(x; c)$ in InfoGAN is defined on the joint distribution of the prior on c and the decoding distribution $p(c)p(x|c)$.

motivate FactorVAE by augmenting the VAE objective with a term that directly encourages independence in the code distribution, arriving at the following objective:

$$\frac{1}{N} \sum_{i=1}^N \left[\mathbb{E}_{q(z|x^{(i)})} [\log p(x^{(i)}|z)] - KL(q(z|x^{(i)})||p(z)) \right] - \gamma KL(q(z)||\bar{q}(z)), \quad (2)$$

where $\bar{q}(z) := \prod_{j=1}^d q(z_j)$. Note that this is also a lower bound on the marginal log likelihood $\mathbb{E}_{p_{data}(x)} [\log p(x)]$. $KL(q(z)||\bar{q}(z))$ is known as *Total Correlation* (TC, [Watanabe, 1960](#)), a popular measure of dependence for multiple random variables. In our case this term is intractable since both $q(z)$ and $\bar{q}(z)$ involve mixtures with a large number of components, and the direct Monte Carlo estimate requires a pass through the entire data set for each $q(z)$ evaluation.³ Hence we take an alternative approach for optimizing this term. We start by observing we can sample from $q(z)$ efficiently by first choosing a datapoint $x^{(i)}$ uniformly at random and then sampling from $q(z|x^{(i)})$. We can also sample from $\bar{q}(z)$ by generating d samples from $q(z)$ and then ignoring all but one dimension for each sample. A more efficient alternative involves sampling a batch from $q(z)$ and then randomly permuting across the batch for each latent dimension (see Alg. 1). This is a standard trick used in the independence testing literature ([Arcones & Gine, 1992](#)) and as long as the batch is large enough, the distribution of these samples will closely approximate $\bar{q}(z)$.

Having access to samples from both distributions allows us to minimise their KL divergence using the *density-ratio trick* ([Nguyen et al., 2010](#); [Sugiyama et al., 2012](#)) which involves training a classifier/discriminator to approximate the density ratio that arises in the KL term. Suppose we have a discriminator D (in our case an MLP) that outputs an estimate of the probability $D(z)$ that its input is a sample from $q(z)$ rather than from $\bar{q}(z)$. Then we have

$$TC(z) = KL(q(z)||\bar{q}(z)) = \mathbb{E}_{q(z)} \left[\log \frac{q(z)}{\bar{q}(z)} \right] \approx \mathbb{E}_{q(z)} \left[\log \frac{D(z)}{1 - D(z)} \right]. \quad (3)$$

We train the discriminator and the VAE jointly. In particular, the VAE parameters are updated using the objective in Eqn. (2), with the TC term replaced using the discriminator-based approximation from Eqn. (3). The discriminator is trained to classify between samples from $q(z)$ and $\bar{q}(z)$, thus learning to approximate the density ratio needed for estimating TC. See Alg. 2 for pseudocode of FactorVAE.

It is important to note that low TC is necessary but not sufficient for meaningful disentangling. For example, when

³We have also tried using a batch estimate of $q(z)$, but this did not work. See Appendix D for details.

Algorithm 1 `permute_dims`

Input: $\{z^{(i)} \in \mathbb{R}^d : i = 1, \dots, B\}$
for $j = 1$ **to** d **do**
 $\pi \leftarrow$ random permutation on $\{1, \dots, B\}$
 $(z_j^{(i)})_{i=1}^B \leftarrow (z_j^{(\pi(i))})_{i=1}^B$
end for
Output: $\{z^{(i)} : i = 1, \dots, B\}$

Algorithm 2 FactorVAE

Input: observations $(x^{(i)})_{i=1}^N$, batch size m , latent dimension d , γ , VAE/Discriminator optimisers: g, g_D
 Initialize VAE and discriminator parameters θ, ψ .
repeat
 Randomly select batch $(x^{(i)})_{i \in \mathcal{B}}$ of size m
 Sample $z_\theta^{(i)} \sim q_\theta(z|x^{(i)}) \forall i \in \mathcal{B}$
 $\theta \leftarrow g(\nabla_\theta \frac{1}{m} \sum_{i \in \mathcal{B}} [\log \frac{p_\theta(x^{(i)}, z_\theta^{(i)})}{q_\theta(z_\theta^{(i)}|x^{(i)})} - \gamma \log \frac{D_\psi(z_\theta^{(i)})}{1 - D_\psi(z_\theta^{(i)})}])$
 Randomly select batch $(x^{(i)})_{i \in \mathcal{B}'}$ of size m
 Sample $z_\theta^{(i)} \sim q_\theta(z|x^{(i)})$ for $i \in \mathcal{B}'$
 $(z_{perm}^{(i)})_{i \in \mathcal{B}'} \leftarrow \text{permute_dims}((z_\theta^{(i)})_{i \in \mathcal{B}'})$
 $\psi \leftarrow g_D(\nabla_\psi \frac{1}{2m} [\sum_{i \in \mathcal{B}} \log(D_\psi(z_\theta^{(i)})) + \sum_{i \in \mathcal{B}'} \log(1 - D_\psi(z_{perm}^{(i)})])$
until convergence of objective.

$q(z|x) = p(z)$, TC=0 but z carries no information about the data. Thus having low TC is only meaningful when we can preserve information in the latents, which is why controlling for reconstruction error is important.

In the GAN literature, divergence minimisation is usually done between two distributions over the data space, which is often very high dimensional (e.g. images). As a result, the two distributions often have disjoint support, making training unstable, especially when the discriminator is strong. Hence it is necessary to use tricks to weaken the discriminator such as instance noise ([Sønderby et al., 2016](#)) or to replace the discriminator with a critic, as in Wasserstein GANs ([Arjovsky et al., 2017](#)). In this work, we minimise divergence between two distributions over the latent space (as in e.g. ([Mescheder et al., 2017](#))), which is typically much lower dimensional and the two distributions have overlapping support. We observe that training is stable for sufficiently large batch sizes (e.g. 64 worked well for $d = 10$), allowing us to use a strong discriminator.

4. A New Metric for Disentanglement

The definition of disentanglement we use in this paper, where a change in one dimension of the representation corresponds to a change in exactly one factor of variation, is

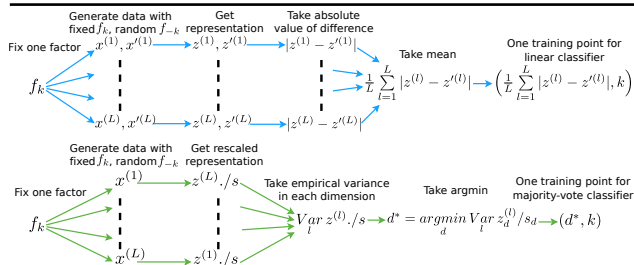


Figure 2. Top: Metric in (Higgins et al., 2016). Bottom: Our new metric, where $s \in \mathbb{R}^d$ is the scale (empirical standard deviation) of latent representations of the full data (or large enough random subset).

clearly a simplistic one. It does not allow correlations among the factors or hierarchies over them. Thus this definition seems more suited to synthetic data with independent factors of variation than to most realistic data sets. However, as we will show below, robust disentanglement is not a fully solved problem even in this simple setting. One obstacle on the way to this first milestone is the absence of a sound quantitative metric for measuring disentanglement.

A popular method of measuring disentanglement is by inspecting *latent traversals*: visualising the change in reconstructions while traversing one dimension of the latent space at a time. Although latent traversals can be a useful indicator of when a model has failed to disentangle, the qualitative nature of this approach makes it unsuitable for comparing algorithms reliably. Doing this would require inspecting a multitude of latent traversals over multiple reference images, random seeds, and points during training. Having a human in the loop to assess the traversals is also too time-consuming and subjective. Unfortunately, for data sets that do not have the ground truth factors of variation available, currently this is the only viable option for assessing disentanglement.

Higgins et al. (2016) proposed a supervised metric that attempts to quantify disentanglement when the ground truth factors of a data set are given. The metric is the error rate of a linear classifier that is trained as follows. Choose a factor k ; generate data with this factor fixed but all other factors varying randomly; obtain their representations (defined to be the mean of $q(z|x)$); take the absolute value of the pairwise differences of these representations. Then the mean of these statistics across the pairs gives one training input for the classifier, and the fixed factor index k is the corresponding training output (see top of Figure 2). So if the representations were perfectly disentangled, we would see zeros in the dimension of the training input that corresponds to the fixed factor of variation, and the classifier would learn to map the index of the zero value to the index of the factor.

However this metric has several weaknesses. Firstly, it could be sensitive to hyperparameters of the linear classifier optimisation, such as the choice of the optimiser and its




Figure 3. A β -VAE model trained on the 2D Shapes data that scores 100% on metric in Higgins et al. (2016) (ignoring the shape factor). First row: originals. Second row: reconstructions. Remaining rows: reconstructions of latent traversals. The model only uses three latent units to capture x -position, y -position, scale and ignores orientation, yet achieves a perfect score on the metric.

hyperparameters, weight initialisation, and the number of training iterations. Secondly, having a linear classifier is not so intuitive – we could get representations where each factor corresponds to a linear combination of dimensions instead of a single dimension. Finally and most importantly, the metric has a failure mode: it gives 100% accuracy even when only $K - 1$ factors out of K have been disentangled; to predict the remaining factor, the classifier simply learns to detect when all the values corresponding to the $K - 1$ factors are non-zero. An example of such a case is shown in Figure 3.

To address these weaknesses, we propose a new disentanglement metric as follows. Choose a factor k ; generate data with this factor fixed but all other factors varying randomly; obtain their representations; normalise each dimension by its empirical standard deviation over the full data (or a large enough random subset); take the empirical variance in each dimension⁴ of these normalised representations. Then the index of the dimension with the lowest variance and the target index k provide one training input/output example for the classifier (see bottom of Figure 2). Thus if the representation is perfectly disentangled, the empirical variance in the dimension corresponding to the fixed factor will be 0. We normalise the representations so that the arg min is invariant to rescaling of the representations in each dimension. Since both inputs and outputs lie in a discrete space, the optimal classifier is the majority-vote classifier (see Appendix B for details), and the metric is the error rate of the classifier. The resulting classifier is a deterministic function of the training data, hence there are no optimisation hyperparameters to tune. We also believe that this metric is conceptually simpler and more natural than the previous one. Most importantly, it circumvents the failure mode of the earlier metric, since the classifier needs to see the lowest variance in a latent dimension for a given factor to classify it correctly.

We think developing a reliable unsupervised disentanglement metric that does not use the ground truth factors is an important direction for future research, since unsupervised

⁴We can use Gini’s definition of variance for discrete latents (Gini, 1971). See Appendix B for details.

disentangling is precisely useful for the scenario where we do not have access to the ground truth factors. With this in mind, we believe that having a reliable supervised metric is still valuable as it can serve as a gold standard for evaluating unsupervised metrics.

5. Related Work

There are several recent works that use a discriminator to optimise a divergence to encourage independence in the latent codes. Adversarial Autoencoder (AAE, [Makhzani et al., 2015](#)) removes the $I(x; z)$ term in the VAE objective and maximizes the negative reconstruction error minus $KL(q(z)||p(z))$ via the density-ratio trick, showing applications in semi-supervised classification and unsupervised clustering. This means that the AAE objective is not a lower bound on the log marginal likelihood. Although optimising a lower bound is not strictly necessary for disentangling, it does ensure that we have a valid generative model; having a generative model with disentangled latents has the benefit of being a single model that can be useful for various tasks e.g. planning for model-based RL, visual concept learning and semi-supervised learning, to name a few. In PixelGAN Autoencoders ([Makhzani & Frey, 2017](#)), the same objective is used to study the decomposition of information between the latent code and the decoder. The authors state that adding noise to the inputs of the encoder is crucial, which suggests that limiting the information that the code contains about the input is essential and that the $I(x; z)$ term should not be dropped from the VAE objective. [Brakel & Bengio \(2017\)](#) also use a discriminator to penalise the Jensen-Shannon Divergence between the distribution of codes and the product of its marginals. However, they use the GAN loss with deterministic encoders and decoders and only explore their technique in the context of Independent Component Analysis source separation.

Early works on unsupervised disentangling include ([Schmidhuber, 1992](#)) which attempts to disentangle codes in an autoencoder by penalising predictability of one latent dimension given the others and ([Desjardins et al., 2012](#)) where a variant of a Boltzmann Machine is used to disentangle two factors of variation in the data. More recently, [Achille & Soatto \(2018\)](#) have used a loss function that penalises TC in the context of supervised learning. They show that their approach can be extended to the VAE setting, but do not perform any experiments on disentangling to support the theory. In a concurrent work, [Kumar et al. \(2018\)](#) used moment matching in VAEs to penalise the covariance between the latent dimensions, but did not constrain the mean or higher moments. We provide the objectives used in these related methods and show experimental results on disentangling performance, including AAE, in Appendix F.

There have been various works that use the notion of pre-

dictability to quantify disentanglement, mostly predicting the value of ground truth factors $f = (f_1, \dots, f_K)$ from the latent code z . This dates back to [Yang & Amari \(1997\)](#) who learn a linear map from representations to factors in the context of linear ICA, and quantify how close this map is to a permutation matrix. More recently [Eastwood & Williams \(2018\)](#) have extended this idea to disentanglement by training a Lasso regressor to map z to f and using its trained weights to quantify disentanglement. Like other regression-based approaches, this one introduces hyperparameters such as the optimiser and the Lasso penalty coefficient. The metric of [Higgins et al. \(2016\)](#) as well as the one we proposed, predict the factor k from the z of images with a fixed f_k but f_{-k} varying randomly. [Schmidhuber \(1992\)](#) quantifies predictability between the different dimensions of z , using a predictor that is trained to predict z_j from z_{-j} .

Invariance and *equivariance* are frequently considered to be desirable properties of representations in the literature ([Goodfellow et al., 2009](#); [Kivinen & Williams, 2011](#); [Lenc & Vedaldi, 2015](#)). A representation is said to be *invariant* for a particular task if it does not change when nuisance factors of the data, that are irrelevant to the task, are changed. An *equivariant* representation changes in a stable and predictable manner when altering a factor of variation. A disentangled representation, in the sense used in the paper, is equivariant, since changing one factor of variation will change one dimension of a disentangled representation in a predictable manner. Given a task, it will be easy to obtain an invariant representation from the disentangled representation by ignoring the dimensions encoding the nuisance factors for the task ([Cohen & Welling, 2014](#)).

Building on a preliminary version of this paper, ([Chen et al., 2018](#)) recently proposed a minibatch-based alternative to our density-ratio-trick-based method for estimating the Total Correlation and introduced an information-theoretic disentangling metric.

6. Experiments

We compare FactorVAE to β -VAE on the following data sets with i) known generative factors: 1) **2D Shapes** ([Matthey et al., 2017](#)): 737,280 binary 64×64 images of 2D shapes with ground truth factors[number of values]: shape[3], scale[6], orientation[40], x-position[32], y-position[32]. 2) **3D Shapes** data: 480,000 RGB $64 \times 64 \times 3$ images of 3D shapes with ground truth factors: shape[4], scale[8], orientation[15], floor colour[10], wall colour[10], object colour[10] ii) unknown generative factors: 3) **3D Faces** ([Paysan et al., 2009](#)): 239,840 grey-scale 64×64 images of 3D Faces. 4) **3D Chairs** ([Aubry et al., 2014](#)): 86,366 RGB $64 \times 64 \times 3$ images of chair CAD models. 5) **CelebA** (cropped version) ([Liu et al., 2015](#)): 202,599 RGB $64 \times 64 \times 3$ images of celebrity faces. The experimental details such as en-

coder/decoder architectures and hyperparameter settings are in Appendix A. The details of the disentanglement metrics, along with a sensitivity analysis with respect to their hyperparameters, are given in Appendix B.

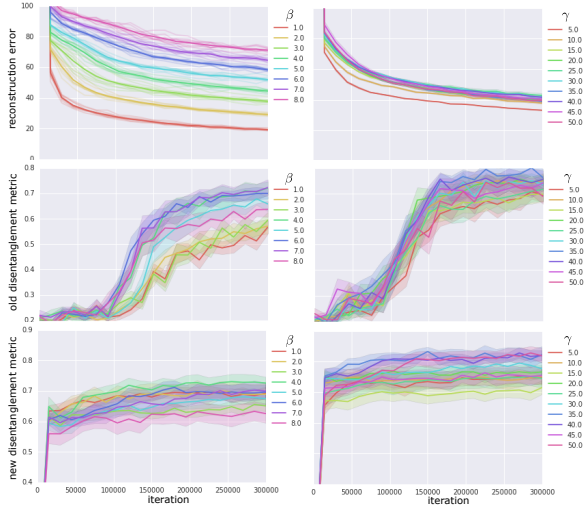


Figure 4. Reconstruction error (top), metric in Higgins et al. (2016) (middle), our metric (bottom). β -VAE (left), FactorVAE (right). The colours correspond to different values of β and γ respectively, and confidence intervals are over 10 random seeds.

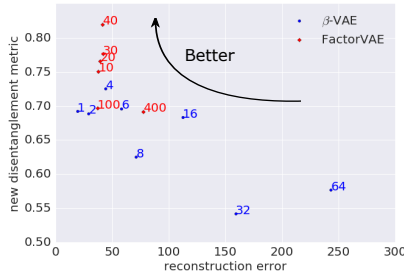


Figure 5. Reconstruction error plotted against our disentanglement metric, both averaged over 10 random seeds at the end of training. The numbers at each point are values of β and γ . Note that we want low reconstruction error and a high disentanglement metric.

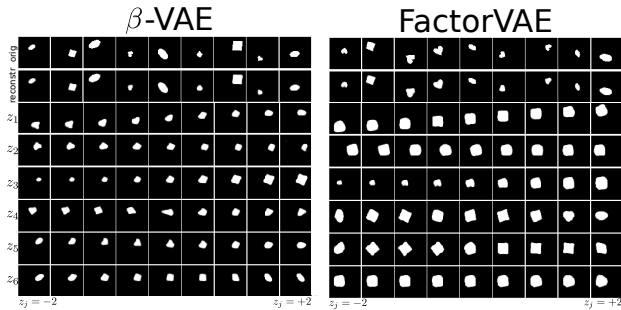


Figure 6. First row: originals. Second row: reconstructions. Remaining rows: reconstructions of latent traversals across each latent dimension sorted by $KL(q(z_j|x)||p(z_j))$, for the best scoring models on our disentanglement metric. Left: β -VAE, score: 0.814, $\beta = 4$. Right: FactorVAE, score: 0.889, $\gamma = 35$.

From Figure 4, we see that FactorVAE gives much better

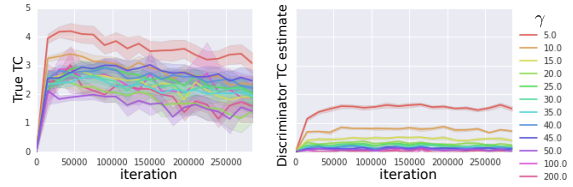


Figure 7. Total Correlation values for FactorVAE on 2D Shapes. Left: True TC value. Right: Discriminator’s estimate of TC.

disentanglement scores than VAEs ($\beta = 1$), while barely sacrificing reconstruction error, highlighting the disentangling effect of adding the Total Correlation penalty to the VAE objective. The best disentanglement scores for FactorVAE are noticeably better than those for β -VAE given the same reconstruction error. This can be seen more clearly in Figure 5 where the best mean disentanglement of FactorVAE ($\gamma = 40$) is around 0.82, significantly higher than the one for β -VAE ($\beta = 4$), which is around 0.73, both with reconstruction error around 45. From Figure 6, we can see that both models are capable of finding x -position, y -position, and scale, but struggle to disentangle orientation and shape, β -VAE especially. For this data set, neither method can robustly capture shape, the discrete factor of variation⁵.

As a sanity check, we also evaluated the correlation between our metric and the metric in Higgins et al. (2016): **Pearson** (linear correlation coefficient): 0.404, **Kendall** (proportion of pairs that have the same ordering): 0.310, **Spearman** (linear correlation of the rankings): 0.444, all with p-value 0.000. Hence the two metrics show a fairly high positive correlation as expected.

We have also examined how the discriminator’s estimate of the Total Correlation (TC) behaves and the effect of γ on the true TC. From Figure 7, observe that the discriminator is consistently underestimating the true TC, also confirmed in (Rosca et al., 2018). However the true TC decreases throughout training, and a higher γ leads to lower TC, so the gradients obtained using the discriminator are sufficient for encouraging independence in the code distribution.

We then evaluated InfoWGAN-GP, the counterpart of InfoGAN that uses Wasserstein distance and gradient penalty. See Appendix G for an overview. One advantage of InfoGAN is that the Monte Carlo estimate of its objective is differentiable with respect to its parameters even for discrete codes c , which makes gradient-based optimisation straightforward. In contrast, VAE-based methods that rely on the reparameterisation trick for gradient-based optimisation require z to be a reparameterisable continuous random variable and alternative approaches require various vari-

⁵This is partly due to the fact that learning discrete factors would require using discrete latent variables instead of Gaussians, but jointly modelling discrete and continuous factors of variation is a non-trivial problem that needs further research.

ance reduction techniques for gradient estimation (Mnih & Rezende, 2016; Maddison et al., 2017). Thus we might expect Info(W)GAN(-GP) to show better disentangling in cases where some factors are discrete. Hence we use 4 continuous latents (one for each continuous factor) and one categorical latent of 3 categories (one for each shape). We tuned for λ , the weight of the mutual information term in Info(W)GAN(-GP), $\in \{0.0, 0.1, 0.2, \dots, 1.0\}$, number of noise variables $\in \{5, 10, 20, 40, 80, 160\}$ and the learning rates of the generator $\in \{10^{-3}, 10^{-4}\}$, discriminator $\in \{10^{-4}, 10^{-5}\}$.

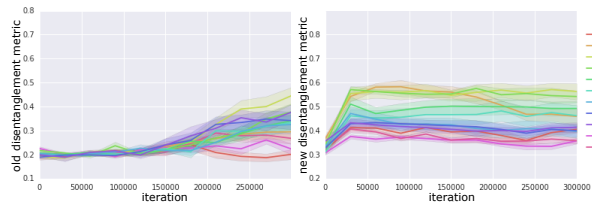


Figure 8. Disentanglement scores for InfoWGAN-GP on 2D Shapes for 10 random seeds per hyperparameter setting. Left: Metric in Higgins et al. (2016). Right: Our metric.




Figure 9. Latent traversals for InfoWGAN-GP on 2D Shapes across four continuous codes (first four rows) and categorical code (last row) for run with best disentanglement score ($\lambda = 0.2$).

However from Figure 8 we can see that the disentanglement scores are disappointingly low. From the latent traversals in Figure 9, we can see that the model learns only the scale factor, and tries to put positional information in the discrete latent code, which is one reason for the low disentanglement score. Using 5 continuous codes and no categorical codes did not improve the disentanglement scores however. InfoGAN with early stopping (before training instability occurs – see Appendix H) also gave similar results. The fact that some latent traversals give blank reconstructions indicates that the model does not generalise well to all parts of the domain of $p(z)$.

One reason InfoWGAN-GP’s poor performance on this data set could be that InfoGAN is sensitive to the generator and discriminator architecture, which is one thing we did not tune extensively. We use a similar architecture to the VAE-based approaches for 2D shapes for a fair comparison, but have also tried a bigger architecture which gave similar results (see Appendix H). If architecture search is indeed important, this would be a weakness of InfoGAN relative to

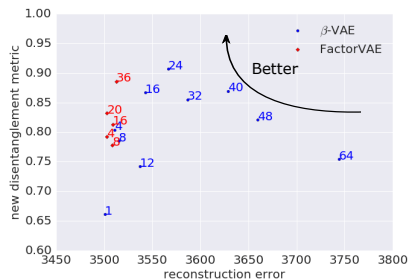


Figure 10. Same as Figure 5 for 3D Shapes data.

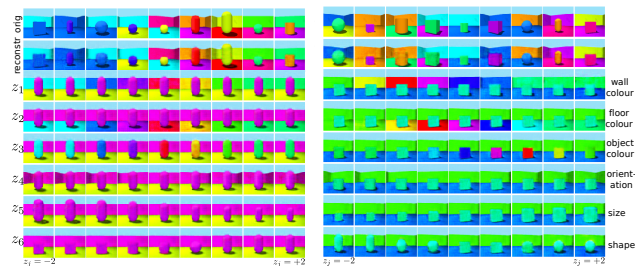


Figure 11. Same as Figure 6 but for 3D Shapes data. Left: β -VAE, score: 1.00, $\beta = 32$. Right: FactorVAE, score: 1.00, $\gamma = 7$.

FactorVAE and β -VAE, which are both much more robust to architecture choice. In Appendix H, we check that we can replicate the results of Chen et al. (2016) on MNIST using InfoWGAN-GP, verify that it makes training stable compared to InfoGAN, and give implementation details with further empirical studies of InfoGAN and InfoWGAN-GP.

We now show results on the 3D Shapes data, which is a more complex data set of 3D scenes with additional features such as shadows and background (sky). We train both β -VAE and FactorVAE for 1M iterations. Figure 10 again shows that FactorVAE achieves much better disentanglement with barely any increase in reconstruction error compared to VAE. Moreover, while the top mean disentanglement scores for FactorVAE and β -VAE are similar, the reconstruction error is lower for FactorVAE: 3515 ($\gamma = 36$) as compared to 3570 ($\beta = 24$). The latent traversals in Figure 11 show that both models are able to capture the factors of variation in the best-case scenario. Looking at latent traversals across many random seeds, however, makes it evident that both models struggled to disentangle the factors for shape and scale.

To show that FactorVAE also gives a valid generative model for both 2D Shapes and 3D Shapes, we present the log marginal likelihood evaluated on the entire data set together with samples from the generative model in Appendix E.

We also show results for β -VAE and FactorVAE experiments on the data sets with unknown generative factors, namely 3D Chairs, 3D Faces, and CelebA. Note that inspecting latent traversals is the only evaluation method possible here. We can see from Figure 12 (and Figures 38 and 39 in Appendix I) that FactorVAE has smaller reconstruction error compared to β -VAE, and is capable of learning sensible

factors of variation, as shown in the latent traversals in Figures 13, 14 and 15. Unfortunately, as explained in Section 4, latent traversals tell us little about the robustness of our method.

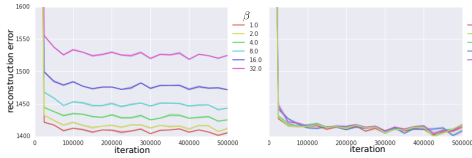


Figure 12. Plots of reconstruction error of β -VAE (left) and FactorVAE (right) for different values of β and γ on 3D Faces data over 5 random seeds.

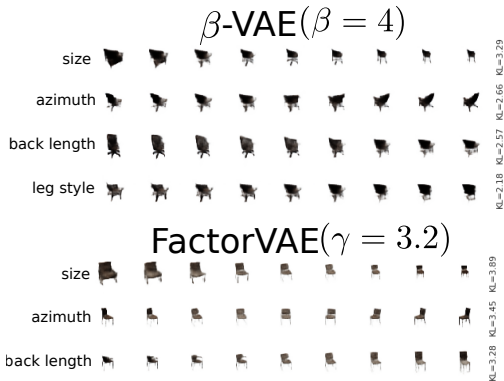


Figure 13. β -VAE and FactorVAE latent traversals across each latent dimension sorted by KL on 3D Chairs, with annotations of the factor of variation corresponding to each latent unit.

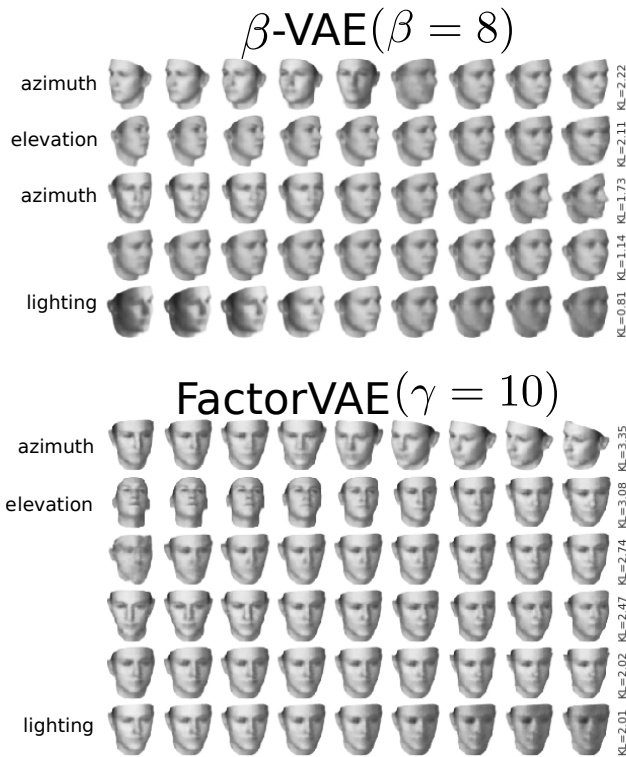


Figure 14. Same as Figure 13 but for 3D Faces.

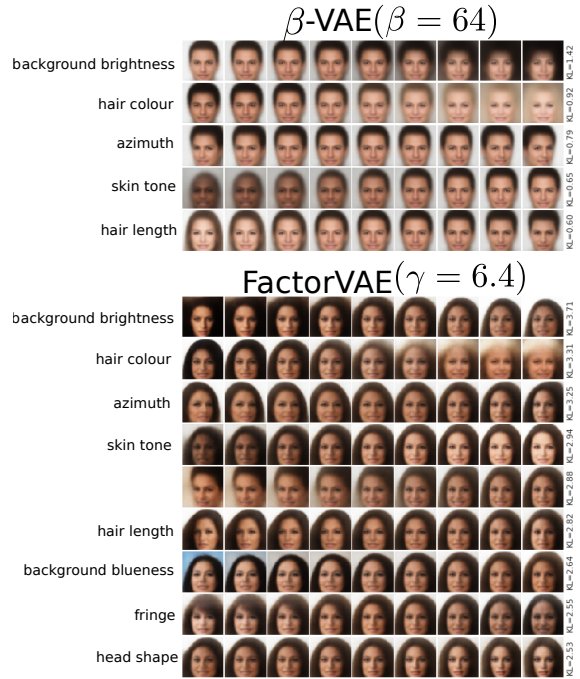


Figure 15. Same as Figure 13 but for CelebA.

7. Conclusion and Discussion

We have introduced FactorVAE, a novel method for disentangling that achieves better disentanglement scores than β -VAE on the 2D Shapes and 3D Shapes data sets for the same reconstruction quality. Moreover, we have identified weaknesses of the commonly used disentanglement metric of Higgins et al. (2016), and proposed an alternative metric that is conceptually simpler, is free of hyperparameters, and avoids the failure mode of the former. Finally, we have performed an experimental evaluation of disentangling for the VAE-based methods and InfoWGAN-GP, a more stable variant of InfoGAN, and identified its weaknesses relative to the VAE-based methods.

One of the limitations of our approach is that low Total Correlation is necessary but not sufficient for disentangling of independent factors of variation. For example, if all but one of the latent dimensions were to collapse to the prior, the TC would be 0 but the representation would not be disentangled. Our disentanglement metric also requires us to be able to generate samples holding one factor fixed, which may not always be possible, for example when our training set does not cover all possible combinations of factors. The metric is also unsuitable for data with non-independent factors of variation.

For future work, we would like to use discrete latent variables to model discrete factors of variation and investigate how to reliably capture combinations of discrete and continuous factors using discrete and continuous latents.

Acknowledgements

We thank Chris Burgess and Nick Watters for providing the data sets and helping to set them up, and thank Guillaume Desjardins, Sergey Bartunov, Mihaela Rosca, Irina Higgins and Yee Whye Teh for helpful discussions.

References

- Achille, A. and Soatto, S. Information Dropout: Learning optimal representations through noisy computation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.
- Arcones, M. A. and Gine, E. On the bootstrap of U and V statistics. *The Annals of Statistics*, pp. 655–674, 1992.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein Generative Adversarial Networks. In *ICML*, 2017.
- Aubry, M., Maturana, D., Efros, A. A., Russell, B. C., and Sivic, J. Seeing 3D chairs: exemplar part-based 2D-3D alignment using a large dataset of cad models. In *CVPR*, 2014.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- Brakel, P. and Bengio, Y. Learning independent features with adversarial nets for non-linear ICA. *arXiv preprint arXiv:1710.05050*, 2017.
- Chen, T. Q., Li, X., Grosse, R., and Duvenaud, D. Isolating sources of disentanglement in variational autoencoders. *arXiv preprint arXiv:1802.04942*, 2018.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. InfoGAN: Interpretable representation learning by information maximizing Generative Adversarial Nets. In *NIPS*, 2016.
- Cohen, T. and Welling, M. Learning the irreducible representations of commutative lie groups. In *ICML*, 2014.
- Denton, E. L. and Birodkar, V. Unsupervised learning of disentangled representations from video. In *NIPS*, 2017.
- Desjardins, G., Courville, A., and Bengio, Y. Disentangling factors of variation via generative entangling. *arXiv preprint arXiv:1210.5474*, 2012.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 12(Jul):2121–2159, 2011.
- Eastwood, C. and Williams, C. A framework for the quantitative evaluation of disentangled representations. In *ICLR*, 2018.
- Gini, C. W. Variability and mutability, contribution to the study of statistical distributions and relations. *Journal of American Statistical Association*, 66:534–544, 1971.
- Goodfellow, I., Lee, H., Le, Q. V., Saxe, A., and Ng, A. Y. Measuring invariances in deep networks. In *NIPS*, 2009.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative Adversarial Nets. In *NIPS*, 2014.
- Goroshin, R., Bruna, J., Tompson, J., Eigen, D., and LeCun, Y. Unsupervised learning of spatiotemporally coherent metrics. In *ICCV*, 2015.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of wasserstein GANs. In *NIPS*, 2017.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. Beta-VAE: Learning basic visual concepts with a constrained variational framework. 2016.
- Hinton, G. E., Krizhevsky, A., and Wang, S. D. Transforming auto-encoders. In *International Conference on Artificial Neural Networks*, pp. 44–51. Springer, 2011.
- Hoffman, M. D. and Johnson, M. J. ELBO surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Advances in Approximate Bayesian Inference, NIPS*, 2016.
- Hsu, W. N., Zhang, Y., and Glass, J. Unsupervised learning of disentangled and interpretable representations from sequential data. In *NIPS*, 2017.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. 2014.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *NIPS*, 2014.
- Kivinen, J. J. and Williams, C. Transformation equivariant boltzmann machines. In *International Conference on Artificial Neural Networks*, 2011.

- Kulkarni, T., Whitney, W. F., Kohli, P., and Tenenbaum, J. Deep convolutional inverse graphics network. In *NIPS*, 2015.
- Kumar, A., Sattigeri, P., and Balakrishnan, A. Variational inference of disentangled latent concepts from unlabeled observations. In *ICLR*, 2018.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building machines that learn and think like people. *Behavioral and Brain Sciences*, pp. 1–101, 2016.
- Lenc, K. and Vedaldi, A. Understanding image representations by measuring their equivariance and equivalence. In *CVPR*, 2015.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3730–3738, 2015.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The CONCRETE distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017.
- Makhzani, A. and Frey, B. PixelGAN autoencoders. In *NIPS*, 2017.
- Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., and Frey, B. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Mathieu, M. F., Zhao, J. J., Ramesh, A., Sprechmann, P., and LeCun, Y. Disentangling factors of variation in deep representation using adversarial training. In *NIPS*, 2016.
- Matthey, L., Higgins, I., Hassabis, D., and Lerchner, A. dSprites: Disentanglement testing Sprites dataset. <https://github.com/deepmind/dsprites-dataset/>, 2017.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial variational Bayes: Unifying Variational Autoencoders and Generative Adversarial Networks. In *ICML*, 2017.
- Mnih, A. and Rezende, D. J. Variational inference for Monte Carlo objectives. In *ICML*, 2016.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 2010.
- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. In *NIPS*, 2016.
- Paysan, P., Knothe, R., Amberg, B., Romdhani, S., and Vetter, T. A 3D face model for pose and illumination invariant face recognition. In *Proceedings of the IEEE International Conference on Advanced Video and Signal based Surveillance*, pp. 296–301, 2009.
- Perry, G., Rolls, E. T., and Stringer, S. M. Continuous transformation learning of translation invariant representations. *Experimental Brain Research*, 204(2):255–270, 2010.
- Reed, S., Sohn, K., Zhang, Y., and Lee, H. Learning to disentangle factors of variation with manifold interaction. In *ICML*, 2014.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.
- Rosca, M., Lakshminarayanan, B., and Mohamed, S. Distribution matching in variational inference. *arXiv preprint arXiv:1802.06847*, 2018.
- Schmidhuber, J. Learning factorial codes by predictability minimization. *Neural Computation*, 4(6):863–879, 1992.
- Siddharth, N., Paige, B., Van de Meent, J. W., Desmaison, A., Wood, F., Goodman, N. D., Kohli, P., and Torr, P. H. S. Learning disentangled representations with semi-supervised deep generative models. In *NIPS*, 2017.
- Sønderby, C. K., Caballero, J., Theis, L., Shi, W., and Huszár, F. Amortised MAP inference for image super-resolution. In *ICLR*, 2016.
- Sugiyama, M., Suzuki, T., and Kanamori, T. Density-ratio matching under the Bregman divergence: a unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64(5):1009–1044, 2012.
- Watanabe, S. Information theoretical analysis of multivariate correlation. *IBM Journal of research and development*, 4(1):66–82, 1960.
- Yang, H. H. and Amari, S. I. Adaptive online learning algorithms for blind separation: maximum entropy and minimum mutual information. *Neural computation*, 9(7): 1457–1482, 1997.