
An Identifiable Double VAE For Disentangled Representations

Graziano Mita^{1,2} Maurizio Filippone¹ Pietro Michiardi¹

Abstract

A large part of the literature on learning disentangled representations focuses on variational autoencoders (VAEs). Recent developments demonstrate that disentanglement cannot be obtained in a fully unsupervised setting without inductive biases on models and data. However, Khemakhem et al., AISTATS, 2020 suggest that employing a particular form of factorized prior, conditionally dependent on auxiliary variables complementing input observations, can be one such bias, resulting in an identifiable model with guarantees on disentanglement. Working along this line, we propose a novel VAE-based generative model with theoretical guarantees on identifiability. We obtain our conditional prior over the latents by learning an optimal representation, which imposes an additional strength on their regularization. We also extend our method to semi-supervised settings. Experimental results indicate superior performance with respect to state-of-the-art approaches, according to several established metrics proposed in the literature on disentanglement.

1. Introduction

Representation learning aims at learning data representations such that it is easier to extract useful information when building classifiers or other predictive tasks (Bengio et al., 2013). Representation learning seeks to obtain the following properties: i) *expressiveness*: a reasonably-sized representation should allow to distinguish among a high number of different input configurations; ii) *abstractness*: learned representations should capture high-level features; iii) *invariance*: representation should be invariant to local changes of input configurations; iv) *interpretability*: learned representations should allow each dimension to be informative about the

¹EURECOM, 06410 Biot (France) ²SAP Labs France, 06250 Mougins (France). Correspondence to: Graziano Mita <graziano.mita@eurecom.fr>, Maurizio Filippone <maurizio.filippone@eurecom.fr>, Pietro Michiardi <pietro.michiardi@eurecom.fr>.

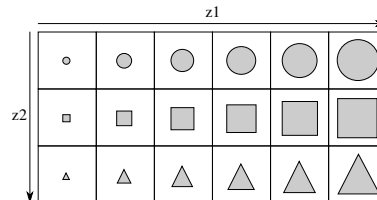


Figure 1. Toy example where each z_i controls a given ground-truth factor: z_1 the size, z_2 the shape of the 2D objects.

given task. These properties are at the core of *disentangled representations*. In disentangled representation learning, the main assumption is that high-dimensional observations \mathbf{x} are the result of a (possibly nonlinear) transformation applied to a low dimensional latent variable of independent generative factors, called *ground-truth factors*, capturing semantically meaningful concepts. Input observations can be thought of as the result of a probabilistic generative process, where latent variables \mathbf{z} are first sampled from a prior distribution $p(\mathbf{z})$, and then the observations \mathbf{x} are sampled from $p(\mathbf{x}|\mathbf{z})$. The goal is to learn a representation of the data that captures the generative factors. In simple terms, illustrated in fig. 1, each dimension of a disentangled representation refers to a single factor of variation.

In this work, we focus on deep generative models, and in particular those based on variational autoencoders (VAE), to learn disentangled representations. A well known theoretical result asserts that disentanglement is essentially impossible in a fully unsupervised setting, without inductive biases on models and data (Locatello et al., 2019). However, inducing a disentangled structure into the latent space where \mathbf{z} lies is feasible by incorporating auxiliary information \mathbf{u} about the ground-truth factors in the model. The type and amount of supervision define different families of disentanglement methods, often classified as supervised, semi-supervised, and weakly-supervised. In most of these methods, the auxiliary variables \mathbf{u} become an integral part of the latent space. However, recent work (Khemakhem et al., 2020) indicates that there are alternative strategies to benefit from auxiliary information, such as using it to impose a structure on the latent space. In their proposal, this is done by learning a prior distribution on the latent space, where the crucial aspect is that this is conditioned on auxiliary information \mathbf{u} that is coupled with every input

observations. Under mild assumptions, it is possible to show that such form of conditioning implies model identifiability, allowing one to recover the original ground-truth factors and therefore providing principled disentanglement.

In this work, we propose a novel generative model that, like Khemakhem et al. (2020), uses a conditional prior and has theoretical identifiability guarantees. We show that our method naturally imposes an optimality constraint, in information theoretic terms, on the conditional prior: this improves the regularization on the function that maps input observations to latent variables, which translates in tangible improvements of disentanglement in practice. Since assuming to have access to auxiliary variables for each input observations, both at training and testing time, is not practical in many applications, we also propose a semi-supervised variant of our method.

Our Contributions: i) We present a detailed overview of VAE-based disentanglement methods using a unified notation. Our focus is on the role of the regularization term. We introduce a distinction between *direct matching approaches*, in which ground-truth factors are directly matched to the latent space, and *indirect matching approaches*, where a prior distribution over the latents is used to structure the learned latent space. ii) We design a new method, that we call *Identifiable Double VAE* (IDVAE) since its ELBO can be seen as a combination of two variational autoencoders, that is identifiable, in theory, and that learns an optimal conditional prior, which is truly desirable in practice. We additionally propose a semi-supervised version of IDVAE to make our method applicable also when auxiliary information is available for a subset of the input observations only. iii) We design an experimental protocol that uses four well-known datasets, and established disentanglement metrics. We compare our method to several state-of-the-art competitors and demonstrate that IDVAE achieves superior disentanglement performance across most experiments.

2. Preliminaries

Let $\mathbf{x} \in \mathbb{R}^n$ be some input observations, which are the result of a transformation of independent latent ground-truth factors $\mathbf{z} \in \mathbb{R}^d$ through a function $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^n$. Then, we have that $\mathbf{x} = \mathbf{f}(\mathbf{z}) + \epsilon$, where ϵ is a Gaussian noise term: $\epsilon \sim \mathcal{N}(0, \Sigma)$, and independent of \mathbf{z} . Let consider the following generative model:

$$p_{\theta}(\mathbf{x}, \mathbf{z}) = p_{\theta}(\mathbf{x}|\mathbf{z})p_{\theta}(\mathbf{z}), \quad (1)$$

where $\theta \in \Theta$ is a vector of model parameters, $p_{\theta}(\mathbf{z}) = \prod_{i=1}^d p_{\theta}(z_i)$ represents the factorized prior probability distribution over the latents and $p_{\theta}(\mathbf{x}|\mathbf{z})$ is the conditional distribution to recover \mathbf{x} from \mathbf{z} . The decoder function $\mathbf{f}(\mathbf{z})$ determines the way \mathbf{z} is transformed into \mathbf{x} within $p_{\theta}(\mathbf{x}|\mathbf{z})$.

Assume to observe some data $\mathcal{D} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$ generated by $p_{\theta^*}(\mathbf{x}, \mathbf{z}) = p_{\theta^*}(\mathbf{x}|\mathbf{z})p_{\theta^*}(\mathbf{z})$, where θ^* are the true, but unknown parameters. Then, the goal is to learn $\theta \in \Theta$ such that:

$$p_{\theta}(\mathbf{x}, \mathbf{z}) = p_{\theta^*}(\mathbf{x}, \mathbf{z}). \quad (2)$$

When eq. (2) holds, it is then possible to recover the generative ground-truth factors. Unfortunately, by observing \mathbf{x} alone, we can estimate the marginal density $p_{\theta}(\mathbf{x}) \approx p_{\theta^*}(\mathbf{x})$, but there are no guarantees about learning the true generative model $p_{\theta^*}(\mathbf{x}, \mathbf{z})$. This is only feasible for models satisfying the following implication:

$$\forall(\theta, \theta') : p_{\theta}(\mathbf{x}) = p_{\theta'}(\mathbf{x}) \implies \theta = \theta'. \quad (3)$$

When eq. (3) holds, the estimated and the true marginal distribution match, and their parameters match too. Then, the model is **identifiable** (Khemakhem et al., 2020) and, as a consequence, it allows one to recover the latent ground-truth factors and obtain a disentangled representation:

$$p_{\theta}(\mathbf{x}) = p_{\theta'}(\mathbf{x}) \implies p_{\theta}(\mathbf{x}, \mathbf{z}) = p_{\theta'}(\mathbf{x}, \mathbf{z}). \quad (4)$$

A practical goal is to aim for model identifiability *up to trivial transformations*, such as permutation and scaling; as long as ground-truth factors can be identified, their order and scale is irrelevant.

3. Related work

Today, a large body of work to learn disentangled representations is based on generative models. In this work, we focus on VAE-based approaches (Kingma & Welling, 2014; Rezende et al., 2014).

Variational Autoencoder. A standard VAE learns the parameters of eq. (1) by introducing an inference model $q_{\phi}(\mathbf{z}|\mathbf{x})$ to derive an ELBO as follows:

$$\mathcal{L}_{\text{VAE}} = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} [\log p_{\theta}(\mathbf{x}|\mathbf{z})] - \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x}) || p(\mathbf{z})), \quad (5)$$

where, by abuse of notation, we write \mathbf{x} in place of $\mathbf{x}^{(i)}$. This avoids clutter in the presentation of VAE-based models, but, clearly, the marginal log-likelihood is composed of a sum of such ELBO terms, one for each observation $\mathbf{x}^{(i)}$ (Kingma & Welling, 2014).

The distribution $p_{\theta}(\mathbf{x}|\mathbf{z})$ has the role of a decoder, whereas $q_{\phi}(\mathbf{z}|\mathbf{x})$ can be seen as an encoder, and it is generally assumed to be a factorized Gaussian with a diagonal covariance matrix. Both distributions are parameterized with neural networks, with parameters θ and variational parameters ϕ . The prior $p(\mathbf{z})$ is generally a factorized, isotropic unit Gaussian.

The first term of eq. (5) relates to the *reconstruction* of the input data using latent variables sampled from the variational approximation of the true posterior. The second term is a *regularization* term, which pushes the approximate posterior $q_\phi(\mathbf{z}|\mathbf{x})$ to match the prior on the latent space. Maximizing eq. (5) across observations implies learning the parameters such that the reconstruction performance is high, and the regularization term is small.

Since both terms that appear in the regularization of eq. (5) are factorized Gaussians with diagonal covariance, one way to interpret the individual components z_i of the latent space is to view them as independent white noise Gaussian channels (Burgess et al., 2017). When the KL term is zero, the latent channels z_i have zero capacity: this happens when the approximate posterior $q_\phi(\mathbf{z}|\mathbf{x})$ matches exactly the prior $p_\theta(\mathbf{z})$. In this case, however, the reconstruction term is penalized. To increase $\text{KL}(\cdot) > 0$, it is necessary to decrease the overlap between channels, and reduce their variances.

Unsupervised disentanglement learning. The above understanding of the regularization term is at the basis of many variants of the original VAE model, that strive to increase the pressure on the regularization term, or elements thereof, to achieve disentanglement, without sacrificing reconstruction properties too much. For example Higgins et al. (2017) propose β -VAE, which modifies eq. (5) by introducing a hyper-parameter β to gauge the pressure on the regularization term throughout the learning process:

$$\mathcal{L}_{\beta\text{-VAE}} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \beta \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \quad (6)$$

When $\beta > 1$, the encoder distribution $q_\phi(\mathbf{z}|\mathbf{x})$ is pushed towards the unit Gaussian prior $p(\mathbf{z})$. In light of the discussion above, the strong penalization of the KL term in β -VAE affects the latent channel distribution, by reducing the spread of their means, and increasing their variances.

Many methods build on β -VAE (Burgess et al., 2017; Kim & Mnih, 2018; Kumar et al., 2018; Chen et al., 2018; Zhao et al., 2019), rewriting the ELBO in slightly different ways. A generalization of the KL term decomposition proposed by Hoffman & Johnson (2016); Makhzani & Frey (2017) is the following (Chen et al., 2018):

$$\mathbb{E}_{\mathbf{x}}[\text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z}))] = I(\mathbf{x}; \mathbf{z}) + \text{KL}(q(\mathbf{z})||\prod_j q(z_j)) + \sum_j \text{KL}(q(z_j)||p(z_j))$$

where $q(\mathbf{z})$ is the aggregated posterior and $I(\mathbf{x}; \mathbf{z})$ is the mutual information between \mathbf{x} and \mathbf{z} . Penalizing $I(\mathbf{x}; \mathbf{z})$ can be harmful to reconstruction purposes, but enforcing a factorized aggregated posterior encourages independence across the dimensions of \mathbf{z} , favouring disentanglement. The dimensional independence in the latent space is encouraged by the second term, known as total correlation (TC). The

third term is a further regularization, preventing the aggregate posterior to deviate too much from the factorized prior.

Note that unsupervised VAE-based approaches approximate the data marginal distribution $p_\theta(\mathbf{x})$, but there are no guarantees to recover the true joint probability distribution $p_\theta(\mathbf{x}, \mathbf{z})$, having access to the input observations \mathbf{x} only (Khemakhem et al., 2020). Pushing the model to learn a representation with statistically independent dimensions is not a sufficient condition to obtain full disentanglement. These considerations were recently formalized in the *impossibility result* (Locatello et al., 2019), but they were already known in the nonlinear ICA literature (Comon, 1994; Hyvärinen & Pajunen, 1999).

Auxiliary variables and disentanglement. To overcome the above limitations, a key idea is to incorporate an inductive bias in the model. The choice of the variational family and prior distribution can be one of such bias (Mathieu et al., 2019; Kumar & Poole, 2020). Alternatively, it is possible to rely on additional information about the ground-truth factors, which we indicate as $\mathbf{u} \in \mathbb{R}^m$. When auxiliary observed variables \mathbf{u} are available, they can be used jointly with \mathbf{z} to reconstruct the original input \mathbf{x} . These methods are usually classified under the semi/weakly supervised family. More specifically Shu et al. (2020) identify three forms of weak supervision: *restricted labeling* (Kingma et al., 2014; Cheung et al., 2015; Siddharth et al., 2017; Klys et al., 2018), *match/group pairing* (Bouchacourt et al., 2018; Hosoya, 2019; Locatello et al., 2020a), and *rank pairing* (Chen & Batmanghelich, 2020a;b). In the extreme case, when all ground-truth factors are known for all the input samples, we label them as supervised disentanglement methods.

As for unsupervised counterpart, methods relying on auxiliary observed variables \mathbf{u} differ in how the regularization term(s) are designed. Some approaches use a “supervised” regularization term to directly match \mathbf{z} and the available ground-truth factors \mathbf{u} : we refer to this form of regularization as *direct matching*. An example is what we here call FULLVAE method (Locatello et al., 2020b), which optimizes the following ELBO:

$$\mathcal{L}_{\text{FULLVAE}} = \mathcal{L}_{\beta\text{-VAE}} - \gamma R_s(q_\phi(\mathbf{z}|\mathbf{x}), \mathbf{u}), \quad (7)$$

where $R_s(\cdot)$ is a loss function between the latent and the ground-truth factors (in the original implementation it is a binary cross entropy loss). Other approaches employ a KL divergence term between the posterior and the prior over the latents: we refer to this form of regularization as *indirect matching*. In other words, direct matching methods require explicit knowledge of one or more ground-truth factors, whereas indirect matching can also use weak information about them. Shu et al. (2020) demonstrated that indirect matching methods can enforce some properties in the latent space, leading to what they define as *consistency* and *restrictiveness*. To obtain full disentanglement, a method must

satisfy both properties on all the latent dimensions. A recent work by Khemakhem et al. (2020) establishes a theoretical framework to obtain model identifiability, which is related to disentanglement. They propose a new generative model called IVAE, that learns a disentangled representation using a factorized prior from the exponential family, crucially conditioned on \mathbf{u} . In practical applications, the conditional prior is chosen to be a Gaussian location-scale family, where the mean and variance of each latent dimension z_i are expressed as a function of \mathbf{u} . Then, it is possible to derive the following ELBO for the IVAE model:

$$\begin{aligned} \mathcal{L}_{\text{IVAE}} = & \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})}[\log p_\theta(\mathbf{x}|\mathbf{z})] \\ & - \beta \text{KL}(q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})||p_\theta(\mathbf{z}|\mathbf{u})). \end{aligned} \quad (8)$$

In eq. (8), we recognize the usual structure of a reconstruction, and a regularization term. A remarkable advancement of the IVAE model relates to identifiability properties: next, we present a new approach to learn an identifiable model that leads to disentangled representations, by using an optimal factorized prior, conditionally dependent on auxiliary observed variables. We also extend our method to deal with more realistic semi-supervised settings.

4. IDVAE: Identifiable Double VAE

Let $\mathbf{x} \in \mathbb{R}^n$, and $\mathbf{u} \in \mathbb{R}^m$ be two observed random variables, and $\mathbf{z} \in \mathbb{R}^d$ a low-dimensional latent variable, with $d \leq n$. Then, consider the following generative models:

$$p_\theta(\mathbf{x}, \mathbf{z}|\mathbf{u}) = p_f(\mathbf{x}|\mathbf{z})p_{\mathbf{T},\eta}(\mathbf{z}|\mathbf{u}), \quad (9)$$

$$p_f(\mathbf{x}|\mathbf{z}) = p_\epsilon(\mathbf{x} - \mathbf{f}(\mathbf{z})), \quad (10)$$

$$p_{\mathbf{T},\eta}(\mathbf{z}|\mathbf{u}) = \prod_i h_i(z_i)g_i(\mathbf{u}) \exp[\mathbf{T}_i(z_i)^\top \boldsymbol{\eta}_i(\mathbf{u})], \quad (11)$$

and

$$p_\vartheta(\mathbf{z}, \mathbf{u}) = p_\vartheta(\mathbf{u}|\mathbf{z})p(\mathbf{z}), \quad (12)$$

where $\theta = (\mathbf{f}, \mathbf{T}, \eta)$ and ϑ are model parameters. Equation (9) corresponds to the process of generating \mathbf{x} given the latents \mathbf{z} . Equation (10) implies that $\mathbf{x} = \mathbf{f}(\mathbf{z}) + \epsilon$, with $\epsilon \sim \mathcal{N}(0, \Sigma)$. We approximate the injective function \mathbf{f} with a neural network. Equation (11) is an exponential conditionally factorial distribution (Bishop, 2006), where h_i is the base measure, $g_i(\mathbf{u})$ is the normalizing constant, $\mathbf{T}_i = [T_{i,1}, \dots, T_{i,k}]^\top$ are the sufficient statistics, and $\boldsymbol{\eta}_i(\mathbf{u}) = [\eta_{i,1}, \dots, \eta_{i,k}]^\top$ are the corresponding parameters. The dimension of each sufficient statistic k is fixed. We choose a conditional prior from the exponential family because it has a good degree of flexibility. Equation (12) formalizes the additional process to obtain \mathbf{u} given \mathbf{z} through $p_\vartheta(\mathbf{u}|\mathbf{z})$, where $p(\mathbf{z})$ is a prior over the latents, usually a factorized, isotropic unit Gaussian.

Given a dataset $\mathcal{D} = \{(\mathbf{x}^{(1)}, \mathbf{u}^{(1)}), \dots, (\mathbf{x}^{(N)}, \mathbf{u}^{(N)})\}$ of observations generated according to eqs. (9) to (12), we are

interested in finding a variational bound \mathcal{L} for the marginal data log-likelihood $p(\mathbf{x}, \mathbf{u})$, which we derive as follows:

$$\log p(\mathbf{x}, \mathbf{u}) = \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{u})||p_\theta(\mathbf{z}|\mathbf{x}, \mathbf{u})) + \mathcal{L}(\theta, \phi),$$

where, by abuse of notation, we write \mathbf{x} and \mathbf{u} in place of $\mathbf{x}^{(i)}$ and $\mathbf{u}^{(i)}$, which we do hereafter as well.

Since the KL term is non-negative, we have the following variational lower bound: $\log p(\mathbf{x}, \mathbf{u}) \geq \mathcal{L}(\theta, \phi)$. Now, we can write the ELBO, which resembles that of eq. (8), but includes an additional term:

$$\begin{aligned} \mathcal{L}(\theta, \phi) = & \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})}[\log p_f(\mathbf{x}|\mathbf{z})] \\ & - \beta \text{KL}(q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})||p_{\mathbf{T},\eta}(\mathbf{z}|\mathbf{u})) + \log p(\mathbf{u}), \end{aligned} \quad (13)$$

where we introduce the parameter β to gauge the pressure on the KL term. Next, focusing on the generative model in eq. (12), we derive the following variational lower bound for $\log p(\mathbf{u})$ in eq. (13), $\log p(\mathbf{u}) \geq \mathcal{L}_{\text{prior}}(\vartheta, \psi)$:

$$\begin{aligned} \mathcal{L}_{\text{prior}}(\vartheta, \psi) = & \mathbb{E}_{q_\psi(\mathbf{z}|\mathbf{u})}[\log p_\vartheta(\mathbf{u}|\mathbf{z})] \\ & - \text{KL}(q_\psi(\mathbf{z}|\mathbf{u})||p(\mathbf{z})), \end{aligned} \quad (14)$$

Combining eq. (13) and eq. (14), we obtain:

$$\begin{aligned} \mathcal{L}_{\text{IDVAE}}(\theta, \phi, \vartheta, \psi) = & \underbrace{\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})}[\log p_f(\mathbf{x}|\mathbf{z})] - \beta \text{KL}(q_\phi(\mathbf{z}|\mathbf{x},\mathbf{u})||p_{\mathbf{T},\eta}(\mathbf{z}|\mathbf{u}))}_{\textcircled{1}} \\ & + \underbrace{\mathbb{E}_{q_\psi(\mathbf{z}|\mathbf{u})}[\log p_\vartheta(\mathbf{u}|\mathbf{z})] - \text{KL}(q_\psi(\mathbf{z}|\mathbf{u})||p(\mathbf{z}))}_{\textcircled{2}}. \end{aligned} \quad (15)$$

We refer the reader to appendix A for the full derivation of eqs. (13) to (15). We call our method IDVAE, Identifiable Double VAE, because it can be seen as the combination of two variational autoencoders $\textcircled{1}$ and $\textcircled{2}$, with independent parameters. In principle, when we optimize the ELBO by summing across all datapoints, e.g. using a doubly stochastic approach (Titsias & Lázaro-Gredilla, 2014) and automatic differentiation, we could treat the two parts separately. However, nothing would prevent $p_{\mathbf{T},\eta}(\mathbf{z}|\mathbf{u})$ and $q_\psi(\mathbf{z}|\mathbf{u})$ to converge to different distributions. Thus, we further make the modeling assumption of constraining the conditional prior in $\textcircled{1}$ to be exactly the variational approximation learned in $\textcircled{2}$, which belongs to the exponential family: in our current implementation, they are both Gaussian distributions, hence we can simply do moment-matching. IDVAE is trained by alternating optimization of eq. (15), where we update part $\textcircled{2}$, then use the conditional prior to update the part $\textcircled{1}$, and loop.

4.1. Identifiability properties

Next, we set up notations and definitions for a general theory of identifiability of generative models (Khemakhem et al., 2020), and show that IDVAE, under mild conditions, is identifiable.

Notation. Concerning the exponential conditionally factorial distribution in eq. (11), we denote by $\mathbf{T}(\mathbf{z})$ the vector of concatenated sufficient statistics defined as follows: $\mathbf{T}(\mathbf{z}) = [\mathbf{T}_1(z_1)^\top, \dots, \mathbf{T}_d(z_d)^\top]^\top \in \mathbb{R}^{dk}$. We denote by $\boldsymbol{\eta}(\mathbf{u})$ the vector of its parameters defined as follows: $\boldsymbol{\eta}(\mathbf{u}) = [\boldsymbol{\eta}_1(\mathbf{u})^\top, \dots, \boldsymbol{\eta}_d(\mathbf{u})^\top]^\top \in \mathbb{R}^{dk}$.

Definition 1. Let \sim be an equivalence relation on the parameter space Θ . We say that eq. (1) is \sim -identifiable if $p_{\boldsymbol{\theta}}(\mathbf{x}) = p_{\boldsymbol{\theta}^*}(\mathbf{x}) \implies \boldsymbol{\theta} \sim \boldsymbol{\theta}^*$.

Definition 2. Let \sim be the equivalence relation on Θ defined as follows: $(\mathbf{f}, \mathbf{T}, \boldsymbol{\eta}) \sim (\mathbf{f}', \mathbf{T}', \boldsymbol{\eta}') \Leftrightarrow \exists \mathbf{A}, \mathbf{c} : \mathbf{T}(\mathbf{f}^{-1}(\mathbf{x})) = \mathbf{A}\mathbf{T}'(\mathbf{f}'^{-1}(\mathbf{x})) + \mathbf{c}, \forall \mathbf{x} \in \mathcal{X}$, where \mathbf{A} is a $dk \times dk$ matrix and \mathbf{c} is a vector of dimension dk . If \mathbf{A} is invertible, we denote this relation by \sim_A .

Definition 2 establishes a specific equivalence relation that allows to recover the sufficient statistics of our model up to a linear matrix multiplication.

Theorem 1. (Khemakhem et al., 2020) Assume we observe data sampled from $p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z}|\mathbf{u}) = p_{\mathbf{f}}(\mathbf{x}|\mathbf{z})p_{\mathbf{T}, \boldsymbol{\eta}}(\mathbf{z}|\mathbf{u})$, where $p_{\mathbf{f}}(\mathbf{x}|\mathbf{z})$ as in eq. (10) and $p_{\mathbf{T}, \boldsymbol{\eta}}(\mathbf{z}|\mathbf{u})$ as in eq. (11), with parameters $\boldsymbol{\theta} = (\mathbf{f}, \mathbf{T}, \boldsymbol{\eta})$. Assume the following holds:

- i The set $\{\mathbf{x} \in \mathcal{X} : \phi_{\epsilon}(\mathbf{x}) = 0\}$ has measure zero, where ϕ_{ϵ} is the characteristic function of the density p_{ϵ} defined in $p_{\mathbf{f}}(\mathbf{x}|\mathbf{z}) = p_{\epsilon}(\mathbf{x} - \mathbf{f}(\mathbf{z}))$.
- ii The function \mathbf{f} is injective.
- iii The sufficient statistics $T_{i,j}$ in eq. (11) are differentiable almost everywhere, and linearly independent on any subset of \mathcal{X} of measure greater than zero.
- iv Being k the dimensionality of the sufficient statistics \mathbf{T}_i in eq. (11) and d the dimensionality of \mathbf{z} , there exist $dk + 1$ distinct point $\mathbf{u}^0, \dots, \mathbf{u}^{dk}$ such that the $dk \times dk$ matrix \mathbf{E} defined as follows is invertible:

$$\mathbf{E} = (\boldsymbol{\eta}(\mathbf{u}^1) - \boldsymbol{\eta}(\mathbf{u}^0); \dots; \boldsymbol{\eta}(\mathbf{u}^{dk}) - \boldsymbol{\eta}(\mathbf{u}^0)) \quad (16)$$

Then the parameters $\boldsymbol{\theta} = (\mathbf{f}, \mathbf{T}, \boldsymbol{\eta})$ are \sim_A -identifiable.

Theorem 1 (sketch of the proof in Appendix C) guarantees a general form of identifiability for IDVAE. Under more restrictive conditions on \mathbf{f} and \mathbf{T} , following the same reasoning of Khemakhem et al. (2020), it is also possible to reduce \mathbf{A} to a permutation matrix.

Note that, in practice, all the VAE-based methods we discuss in this work are approximate. When using a simple,

synthetic dataset, where the generative process is controlled, full disentanglement can be verified experimentally (Khemakhem et al., 2020). However, in a realistic setting, the modeling choice for both $q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{u})$ and $q_{\psi}(\mathbf{z}|\mathbf{u})$ can have an impact on disentanglement. Even when recognition models have enough capacity to fit the data (in our experiments they are Gaussian with diagonal covariance), theoretical guarantees might still fall short, despite the availability of auxiliary variables for all input observations. This could be due to, for example, suboptimal solutions found by the optimization algorithm or to the finite data regime.

4.2. Learning an optimal conditional prior

In this paper, we advocate for a particular form of a conditional prior, that is the result of learning an optimal representation \mathbf{z} , of auxiliary, observed variables \mathbf{u} .

In general, an *optimal* representation, for a generic task \mathbf{y} (in our case, we aim at reconstructing \mathbf{u}) is defined in terms of sufficiency and minimality: \mathbf{z} is *sufficient* for the task \mathbf{y} if $I(\mathbf{u}; \mathbf{y}) = I(\mathbf{z}; \mathbf{y})$, where $I(\cdot; \cdot)$ is the mutual information; \mathbf{z} is *minimal* if it compresses the input such that it discards all variability that is not relevant for the task (Achille & Soatto, 2016). As shown in (Tishby et al., 1999), the so called *Information Bottleneck* (IB) can be used to learn an optimal representation \mathbf{z} for the task \mathbf{y} , which amounts to optimizing the following Lagrangian:

$$\mathcal{L}_{\text{IB}} = H(\mathbf{y}|\mathbf{u}) + \beta I(\mathbf{u}; \mathbf{z}), \quad (17)$$

where we denote the entropy by $H(\cdot)$, with the constant β controlling the trade-off between sufficiency and minimality. It is easy to show that eq. (17) and eq. (14) are equivalent (with $\beta = 1$) when the task is reconstruction.

In our method, we learn the conditional prior $q_{\psi}(\mathbf{z}|\mathbf{u})$ in part ② of eq. (15), and use it in part ① by setting $p_{\mathbf{T}, \boldsymbol{\eta}}(\mathbf{z}|\mathbf{u}) = q_{\psi}(\mathbf{z}|\mathbf{u})$. In light of above discussion, this is equivalent to imposing an additional constraint that pushes the conditional prior to learn an *optimal* representation \mathbf{z} from \mathbf{u} ; the KL term of part ① in eq. (15) pushes $q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{u})$ toward the optimal conditional prior, which results in superior regularization quality.

Note that Theorem 1 requires auxiliary variables \mathbf{u} to be expressive enough to recover all the independent factors through the parameters $\boldsymbol{\eta}(\mathbf{u})$. In information theoretic terms, \mathbf{u} must be sufficient to recover the ground-truth factors, but there is no explicit need for the extra optimality constraint on $p_{\mathbf{T}, \boldsymbol{\eta}}(\mathbf{z}|\mathbf{u})$. While Theorem 1 remains valid for an optimal conditional prior, we demonstrate experimentally that, when variational approximations, sub-optimal solutions, or finite data size spoil theoretical results, learning an optimal conditional prior is truly desirable.

4.3. A semi-supervised variant of IDVAE

So far, we worked under the assumption that the auxiliary information \mathbf{u} is consistently available for every \mathbf{x} . In real scenarios, it is more likely to observe \mathbf{u} for a subset of the input observations. Thus, we propose a variation of IDVAE for a semi-supervised setting. We consider a new objective function that consists of two terms (Kingma et al., 2014):

$$\mathcal{L}_{\text{SS-IDVAE}} = \sum_{(\mathbf{x}, \mathbf{u}) \sim p_l} \mathcal{L}_l(\mathbf{x}, \mathbf{u}) + \sum_{\mathbf{x} \sim p_u} \mathcal{L}_u(\mathbf{x}), \quad (18)$$

$$\mathcal{L}_l(\mathbf{x}, \mathbf{u}) = \mathcal{L}_{\text{IDVAE}}(\mathbf{x}, \mathbf{u}), \quad (19)$$

$$\mathcal{L}_u(\mathbf{x}) = \mathbb{E}_{q_\zeta(\mathbf{u}|\mathbf{x})}[\mathcal{L}_l(\mathbf{x}, \mathbf{u})] + \mathcal{H}(q_\zeta(\mathbf{u}|\mathbf{x})), \quad (20)$$

where \mathcal{L}_l and \mathcal{L}_u are the labeled and unlabeled terms respectively; $q_\zeta(\mathbf{u}|\mathbf{x})$ in eq. (20) is used to derive \mathbf{u} from \mathbf{x} when \mathbf{u} is not provided as input. To be precise, we should add to eq. (18) a third term $-\mathbb{E}_{(\mathbf{x}, \mathbf{u}) \sim p_l}[\log q_\zeta(\mathbf{u}|\mathbf{x})]$ – such that it can learn also from labeled data. Clearly, this method also applies to the work from Khemakhem et al. (2020).

5. Experiments

5.1. Experimental settings

Methods. We compare IDVAE against three disentanglement methods: β -VAE, FULLVAE, IVAE. β -VAE (Higgins et al., 2017) is a baseline for indirect matching methods where no ground-truth factor is known at training time and the only way to enforce a disentangled representation is by increasing the strength of the regularization term through the hyper-parameter β . FULLVAE (Locatello et al., 2020b) is the representative of direct matching methods: it can be considered as a standard β -VAE with an additional regularization term, weighted by an hyper-parameter γ , to match the latent space to the target ground-truth factors. As done in the original implementation, we use a binary cross entropy loss for FULLVAE, where the targets are normalized in $[0, 1]$; we also set $\beta = 1$, to measure the impact of the supervised loss term only. IVAE (Khemakhem et al., 2020) is another indirect matching method where the regularization term, weighted again by β , involves a conditional prior. We additionally report the results for the semi-supervised versions of FULLVAE, IVAE, and IDVAE, which we denote as SS-FULLVAE, SS-IVAE¹, SS-IDVAE, respectively. Variational approximations, and the conditional priors, are Gaussian distributions with diagonal covariance. All methods have been implemented in PyTorch (Paszke et al., 2019).

Datasets. We consider four common datasets in the disentanglement literature, where observations are images built as a deterministic function of known generative factors:

¹The original work (Khemakhem et al., 2020) is not semi-supervised. We extended it for our comparative analysis.

Dataset	Size	Ground-truth factors (distinct values)
DSPRITES	737'280	shape(3), scale(6), orientation(40), x(32), y(32)
CARS3D	17'568	elevation(4), azimuth(24), object type (183)
SHAPES3D	480'000	floor color(10), wall color(10), object color(8), object size(8), object type(4), azimuth(15)
SMALLNORB	24'300	category(5), elevation(9), azimuth(18), light(6)

Table 1. Main characteristics of the datasets.

DSPRITES (Higgins et al., 2017), SHAPES3D (Kim & Mnih, 2018), CARS3D (Reed et al., 2015) and SMALLNORB (Le-Cun et al., 2004). We have full control on the generative process and explicit access to the ground-truth factors. All ground-truth factors are normalized in the range $[0, 1]$; for discrete factors, we implicitly assume an ordering before applying normalization. All images are reshaped to a 64×64 size. A short description of the datasets is reported in table 1. Implementations of the generative process for each dataset are based on the code provided by Locatello et al. (2019).

Disentanglement metrics. In the literature, several metrics have been proposed to measure disentanglement, with known advantages and disadvantages, and ability to capture different aspects of disentanglement. We report the results for some of the most popular metrics: beta score (Higgins et al., 2017), MIG (Chen et al., 2018), SAP (Kumar et al., 2018), modularity and explicitness (Ridgeway & Mozer, 2018), all with values between 0 and 1. The implementation of the metrics is based on Locatello et al. (2019). We refer the reader to appendix E for further details.

Experimental protocol. In order to fairly evaluate the impact of the regularization terms, all the tested methods have the same convolutional architecture (widely adopted in most recent works), optimizer, hyper-parameters of the optimizer and batch size. The latent dimension \mathbf{z} is fixed to the true number of ground-truth factors. The conditional prior in IVAE is a MLP network; in IDVAE we use a simple MLP VAE. The same architecture is taken for the conditional prior of the semi-supervised counterparts. Moreover, $q_\zeta(\mathbf{u}|\mathbf{x})$ is implemented by a convolutional neural network. We refer the reader to the appendix D for more details.

We tried six different values of regularization strength associated to the target regularization term of each method – β for β -VAE, IVAE and IDVAE, and γ for FULLVAE: $[1, 2, 4, 6, 8, 16]$. These are recurring values in the disentanglement literature. For each model configuration and dataset, we run the training procedure with 10 random seeds, given that all methods are susceptible to initialization values. After 300'000 training iterations, every model is evaluated according to the disentanglement metrics described above. For FULLVAE, IVAE and IDVAE, all ground-truth factors are provided as input, although IVAE and IDVAE work as well with a subset of them (or with any other addition-

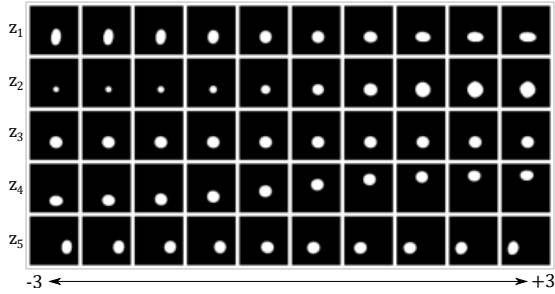


Figure 2. Latent traversal of IDVAE model trained on DSPRITES.

ally observed variable). We apply the same protocol for the semi-supervised experiments too, where we provide, at training time, all the ground-truth factors for a subset of the input observations only, 1% and 10% respectively. At testing time, \mathbf{u} is instead estimated from $q_{\zeta}(\mathbf{u}|\mathbf{x})$.

5.2. Experimental results

Qualitative Evaluation. Latent traversal is a simple approach to visualize disentangled representations, by plotting the effects that each latent dimension of a randomly selected sample has on the reconstructed output. In fig. 2, we evaluate a configuration (single seed) of our IDVAE model trained on DSPRITES (other datasets in appendix F). Every row of the figure represents a latent dimension that we vary in the range $[-3, 3]$, while keeping the other dimensions fixed. We can see that z_1 has learned orientation reasonably well; z_2 is responsible of the object scale; z_4 and z_5 reflect changes on the vertical and horizontal axis, respectively. z_3 tried to learn, without success, shape changes. Next, we rely on disentanglement metrics to make a quantitative comparison among the tested methods.

Disentanglement Evaluation. In fig. 3, we report, for each method and for each dataset, the ranges of the beta score and explicitness values with a box-plot (results in tabular format can be found in appendix F). The variance of the box-plots is due to the random seeds and regularization strengths, which are the only parameters we vary. Furthermore, fig. 3 includes the results for SS-IDVAE, SS-IVAE and SS-FULLVAE (trained with 1% and 10% labeled samples), with different shades of green, blue, and red, respectively. The remaining evaluation metrics can be found in appendix F, but they are essentially all correlated, as also noted in Locatello et al. (2019).

Overall, we observe, as expected, that β -VAE is often the worst method. Indeed, it has no access to any additional information at training time except the data itself. Despite this, β -VAE disentanglement performance is surprisingly not that far from FULLVAE that directly matches the latent space with the ground-truth factors. In some cases, β -VAE obtains very high beta scores (see outliers), such as for

DSPRITES and CARS3D datasets, confirming the sensitivity to random initialization of unsupervised methods (Locatello et al., 2019). Note also that FULLVAE exhibits inconsistent performance across the four datasets.

IDVAE emerges as the best method across several disentanglement metrics, except for SMALLNORB, where FULLVAE’s beta score is slightly better. For this specific dataset and metric, there are no considerable differences among methods, since most of the box-plots overlap. We note that IDVAE outperforms IVAE: considering that the two methods differ for the way the conditional prior is learned, our experiments show that an optimal conditional prior, as we propose in this work, offers substantial benefits in terms of disentanglement and it is the only reason for IDVAE superiority. Finally, although both IVAE and IDVAE have theoretical guarantees on disentanglement and use the full set of ground-truth factors as input, they do not always obtain the maximum evaluation score, in practice. This is in line with the considerations in section 4.1.

The analysis above remains valid if we consider the semi-supervised versions of the tested methods, too. We observe that, with the exception of SMALLNORB, SS-IDVAE’s disentanglement performance coherently increases when it observes more labeled instances. The same trend is generally followed by SS-IVAE. SS-FULLVAE, instead, seems to be less susceptible to the number of labeled instances. In general, even a small percentage of labeled instances (1%) is enough for SS-IDVAE to outperform β -VAE and to keep up with FULLVAE that is, however, a fully supervised method. This suggests that SS-IDVAE is a valid choice for applications where collecting additional information about the training data is difficult or expensive.

Impact of the regularization strength. The disentanglement performance of each method might change drastically as a function of the regularization strength: some approaches might work significantly better in some ranges and very badly in others. In fig. 4, we plot, for each method and for each dataset, the median of the beta score and explicitness evaluation values as a function of the regularization strength. This is also useful to see if there are methods that consistently dominate others. In this case, we do not report the results for SS-IDVAE, SS-IVAE, and SS-FULLVAE to make the plots more easily readable. Additional disentanglement score results, including the semi-supervised versions, can be found in appendix F.

Across all the datasets, IDVAE achieves the best median scores for a wide range of regularization strengths. In DSPRITES, CARS3D, and SHAPES3D, IDVAE dominates all the other methods (IDVAE is largely dominant also considering the remaining evaluations metrics). The performance of IVAE and FULLVAE can match that of ID-

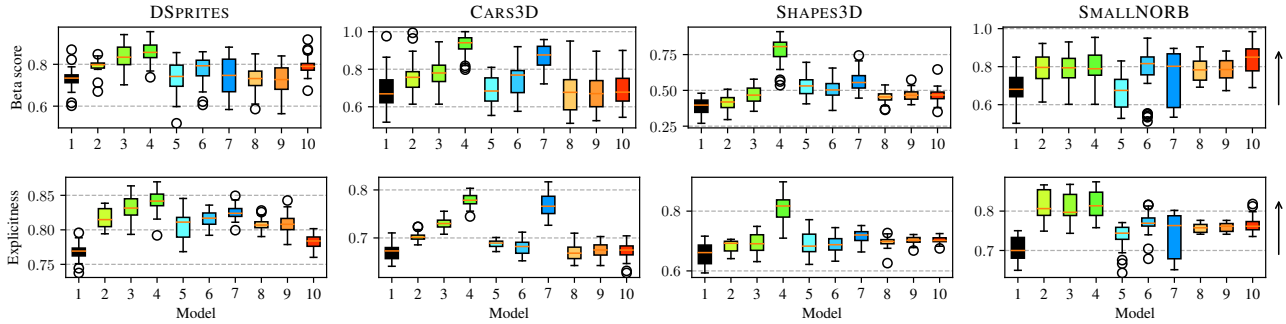


Figure 3. Beta score and explicitness (the higher the better). 1= β -VAE, 2=SS-IDVAE (1%), 3=SS-IDVAE (10%), 4=IDVAE, 5=SS-IDVAE (1%), 6=SS-IDVAE (10%), 7=IVAE, 8=SS-FULLVAE (1%), 9=SS-FULLVAE (10%), 10=FULLVAE. Percentage of labeled samples in parenthesis.

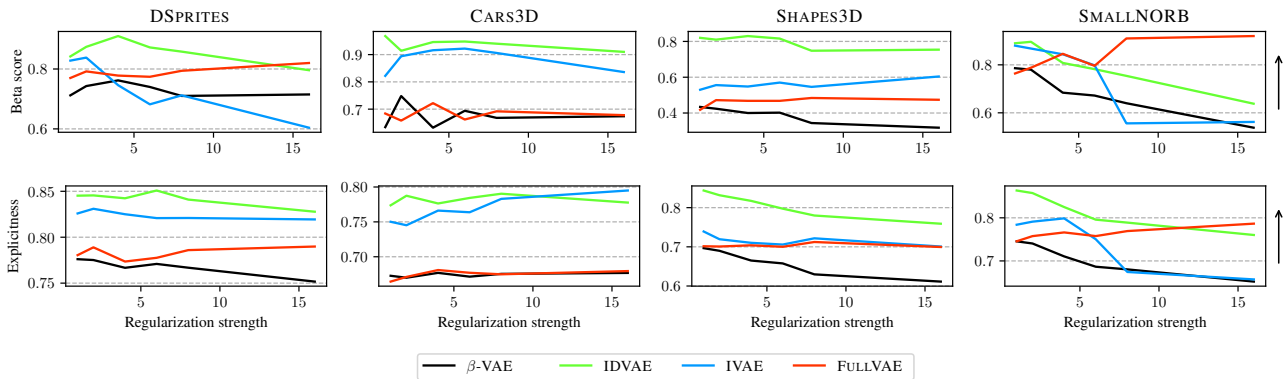


Figure 4. Beta score and explicitness median (the higher the better) as a function of the regularization strength.

VAE in some datasets, but the behavior is not consistent: if we focus on beta score, IVAE is the second best method in CARS3D and SHAPES3D, whereas in DSPRITES and SMALLNORB, performance drops when we increase the regularization strength – even β -VAE performs better; FULLVAE behaves well for DSPRITES and SMALLNORB, but it is on par only with β -VAE in CARS3D and SHAPES3D.

By observing the evolution of the disentanglement scores, it appears that there is no clear strategy to choose the regularization strength. For IDVAE, in datasets such as DSPRITES and CARS3D, the regularization strength does not significantly affect the beta score; in SHAPES3D and SMALLNORB, we note instead a decreasing monotonic trend. The situation is similar if we look at the explicitness, but it differs if we consider different disentanglement metrics. It is plausible to deduce that the regularization strength is both model and data specific, and it is also affected by the choice of the disentanglement metric.

5.3. Limitations

In our experimental campaign we use the same convolutional architecture for all the methods we compare. We do not vary the optimization hyper-parameters and the dimen-

sion of the latent variables. Hence, we cannot ensure that every method runs in its best conditions. Nevertheless, our experimental protocol is in line with the standards in the literature of disentangled representation learning, makes our analysis independent of method-specific optimizations, and has the benefit of reducing training times.

Also, we use the whole set of ground-truth factors as auxiliary variables, in the semi-supervised settings too, whereas it is possible to study the impact of only a subset of the factors to be available. Moreover, IDVAE and IVAE can use any kind of auxiliary variables, as long as they are informative about the ground-truth factors: they are not restricted to using, e.g., labels corresponding to input data, as we (and many other studies) do in our experiments.

Finally, we do not study the implications and benefits of disentanglement for solving complex downstream tasks, which is an interesting task that we leave for future work.

6. Conclusion

In this work, we made a step further in the design of identifiable generative models to learn disentangled representations.

IDVAE uses a prior that encodes ground-truth factor information captured by auxiliary observed variables. The key idea was to learn an optimal representation of the latent space, defined by an inference network on the posterior of the latent variables, given the auxiliary variables. Such posterior is then used as a prior on the latent variables of a second generative model, whose inference network learns a mapping between input observations and latents. We also proposed a semi-supervised version of IDVAE that can be applied when auxiliary variables are available for a subset

of the input observations only. Experimental results offer evidence that IDVAE and SS-IDVAE often outperform existing alternatives to learn disentangled representations, according to several established metrics.

ACKNOWLEDGEMENTS

The authors thank SAP Labs France for their support. Maurizio Filippone gratefully acknowledges support from the AXA Research Fund.

References

- Achille, A. and Soatto, S. Information dropout: learning optimal representations through noise. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2013.
- Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, 2006.
- Bouchacourt, D., Tomioka, R., and Nowozin, S. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. In *Proc. of the 32nd AAAI Conf. on Artif. Intel.*, AAAI, 2018.
- Burgess, C. P., Higgins, I., Pal, A., Matthey, L., Watters, N., Desjardins, G., and Lerchner, A. Understanding disentangling in β -vae. In *Proc. of the 30th Int. Conf. on Neural Inf. Proc. Sys.*, NeurIPS, 2017.
- Chen, J. and Batmanghelich, K. Weakly supervised disentanglement by pairwise similarities. In *Proc. of the 34th AAAI Conf. on Artif. Intel.*, AAAI, 2020a.
- Chen, J. and Batmanghelich, K. Robust ordinal vae: Employing noisy pairwise comparisons for disentanglement. *ArXiv*, 2020b.
- Chen, T. Q., Li, X., Grosse, R. B., and Duvenaud, D. K. Isolating sources of disentanglement in variational autoencoders. In *Proc. of the 31st Int. Conf. on Neural Inf. Proc. Sys.*, NeurIPS, 2018.
- Cheung, B., Livezey, J. A., Bansal, A. K., and Olshausen, B. A. Discovering hidden factors of variation in deep networks. In *CoRR*, 2015.
- Comon, P. Independent component analysis, a new concept? *Signal Process.*, 36:287–314, 1994.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M. M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. In *Proc. of the 5th Int. Conf. on Learn. Repr.*, ICLR, 2017.
- Hoffman, M. D. and Johnson, M. J. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Adv. in Approx. Bayes. Infer.*, NeurIPS, 2016.
- Hosoya, H. Group-based learning of disentangled representations with generalizability for novel contents. In *Proc. of the 28th Int. Joint Conf. on Artif. Intel.*, IJCAI, 2019.
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural networks*, 12:429–439, 1999.
- Khemakhem, I., Kingma, D. P., Mont, R. P., and Hyvärinen, A. Variational autoencoders and nonlinear ica: A unifying framework. In *Proc. of the 23rd Int. Conf. on Artif. Intel. and Stat.*, AISTATS, 2020.
- Kim, H. and Mnih, A. Disentangling by factorising. In *Proc. of the 35th Int. Conf. on Mach. Learn.*, ICML, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *Proc. of the 2nd Int. Conf. on Learn. Repr.*, ICLR, 2014.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In *Proc. of the 27th Int. Conf. on Neural Inf. Proc. Sys.*, NeurIPS, 2014.
- Klys, J., Snell, J., and Zemel, R. Learning latent subspaces in variational autoencoders. In *Proc. of the 31st Int. Conf. on Neural Inf. Proc. Sys.*, NeurIPS, 2018.
- Kumar, A. and Poole, B. On implicit regularization in β -VAEs. In *Proceedings of the 37th International Conference on Machine Learning*, ICML, 2020.
- Kumar, A., Sattigeri, P., and Balakrishnan, A. Variational inference of disentangled latent concepts from unlabeled observations. In *Proc. of the 6th Int. Conf. on Learn. Repr.*, ICLR, 2018.
- LeCun, Y., Huang, F. J., and Bottou, L. Learning methods for generic object recognition with invariance to pose and lighting. In *Proc. of the 2004 IEEE Comput. Society Conf. on Comput. Vision and Pat. Recogn.*, CVPR, 2004.
- Locatello, F., Bauer, S., Lucic, M., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled representations. In *Proc. of the 36th Int. Conf. on Mach. Learn.*, ICML, 2019.
- Locatello, F., Poole, B., Rätsch, G., Schölkopf, B., Bachem, O., and Tschannen, M. Weakly-supervised disentanglement without compromises. In *Proc. of the 37th Int. Conf. on Mach. Learn.*, ICML, 2020a.
- Locatello, F., Tschannen, M., Bauer, S., Rätsch, G., Schölkopf, B., and Bachem, O. Disentangling factors of variation using few labels. In *Proc. of the 8th Int. Conf. on Learn. Repr.*, ICLR, 2020b.
- Makhzani, A. and Frey, B. J. Pixelgan autoencoders. In *Proc. of the 30th Int. Conf. on Neural Inf. Proc. Sys.*, NeurIPS, 2017.

- Mathieu, E., Rainforth, T., Siddharth, N., and Teh, Y. W. Disentangling disentanglement in variational autoencoders. In *Proceedings of the 36th International Conference on Machine Learning, ICML, 2019*.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Proc. of the 32nd Int. Conf. on Neural Inf. Proc. Sys., NeurIPS, 2019*.
- Reed, S. E., Zhang, Y., Zhang, Y., and Lee, H. Deep visual analogy-making. In *Proc. of the 28th Int. Conf. on Neural Inf. Proc. Sys., NeurIPS, Curran Associates, Inc., 2015*.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *Proc. of the 31st Int. Conf. on Mach. Learn., ICML, 2014*.
- Ridgeway, K. and Mozer, M. C. Learning deep disentangled embeddings with the f-statistic loss. In *Proc. of the 31st Int. Conf. on Neural Inf. Proc. Sys., NeurIPS, 2018*.
- Shu, R., Chen, Y., Kumar, A., Ermon, S., and Poole, B. Weakly supervised disentanglement with guarantees. In *Proc. of the 8th Int. Conf. on Learn. Repr., ICLR, 2020*.
- Siddharth, N., Paige, B., van de Meent, J., Desmaison, A., Goodman, N., Kohli, P., Wood, F., and Torr, P. Learning disentangled representations with semi-supervised deep generative models. In *Proc. of the 30th Int. Conf. on Neural Inf. Proc. Sys., NeurIPS, 2017*.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. In *Proc. of the 34th Annual Allert. Conf. on Comm. Contr and Comput., 1999*.
- Titsias, M. and Lázaro-Gredilla, M. Doubly stochastic variational bayes for non-conjugate inference. In *Proc. of the 31st Int. Conf. on Mach. Learn., ICML, 2014*.
- Zhao, S., Song, J., and Ermon, S. Infovae: Balancing learning and inference in variational autoencoders. In *Proc. of the 33rd AAAI Conf. on Artif. Intel., AAAI, 2019*.