Disentangling Domain and Content

Dan Andrei Iliescu* University of Cambridge Cambridge, UK dai24@cam.ac.uk Aliaksei Mikhailiuk University of Cambridge Cambridge, UK am2442@cam.ac.uk

am2442@cam.ac.uk dj Rafal Mantiuk University of Cambridge

Damon Wischik University of Cambridge Cambridge, UK djw1005@cam.ac.uk

Abstract

Cambridge, UK rkm38@cam.ac.uk

Many real-world datasets can be divided into groups according to certain salient features (e.g. grouping images by subject, grouping text by font, etc.). Often, machine learning tasks require that these features be represented separately from those manifesting independently of the grouping. For example, image translation entails changing the style of an image while preserving its content. We formalize these two kinds of attributes as two complementary generative factors called "domain" and "content", and address the problem of disentangling them in a fully unsupervised way. To achieve this, we propose a principled, generalizable probabilistic model inspired by the Variational Autoencoder. Our model exhibits state-of-the-art performance on the composite task of generating images by combining the domain of one input with the content of another. Distinctively, it can perform this task in a few-shot, unsupervised manner, without being provided with explicit labelling for either domain or content. The disentangled representations are learned through the combination of a group-wise encoder and a novel domain-confusion loss.

1 Introduction

Learning rich, interpretable representations with deep neural networks is one of the main challenges of current artificial intelligence research. Achieving such representations would enable us to perform complex and highly useful operations on high-dimensional data (Bengio et al., 2013). Perhaps the first milestone that has yet to be reached in this research is learning representations which easily factorize along the lines of recognizable human concepts. This property is called "the disentanglement of generative factors", and is an accelerating field of inquiry (Tschannen et al., 2018), with many major contributions coming from models based on the Variational Autoencoder (Kingma & Welling, 2014; Rezende et al., 2014).

Recent work (Locatello et al., 2019; van Steenkiste et al., 2019) has revealed limitations in the current methods caused by the inherent ambiguity of the disentanglement objective. They have pointed out the need for equipping models with inductive biases appropriate to their respective application. An example disentanglement task with such increased specificity is the learning to generate multi-object scenes whereby the representation is trained to factorize along object lines (Engelcke et al., 2020; Burgess et al., 2019; Greff et al., 2019).

As a further step towards this goal, we identify another promising disentanglement objective, namely the separation of domain and content representations, widely applicable to a variety of tasks, ranging

Preprint. Under review.

from unsupervised translation to missing data imputation. In a general sense, whenever there exists some form of grouping imposed on a dataset, the notion of domain arises naturally to characterize the attributes of the data which are common within groups but differ across groups. Such attributes could be the style of a painting in the context of style transfer, or the language of a text in the context of neural machine translation. The notion of content then appears as a counterpart to the domain to encompass the features which occur independently of the domain features. For instance, in the context of style transfer, the actual subject of the painting represents the content.

1.1 Related Work

There are many research directions which lead into the domain-content paradigm. Early work on domain adaptation (Ben-David et al., 2010; Ganin et al., 2016), for example, has highlighted the desirability of learning domain-invariant (content) representations of the data, in order to perform classification and regression in a common space. The model of Gonzalez-Garcia et al. (2018) can successfully separate domain-specific from domain-invariant features for two domains.

Problems such as image-to-image translation, which entails changing the domain of an image while preserving its content, have been extensively studied. Major deep learning innovations have come from this area (Isola et al., 2017; Zhu et al., 2017), producing results of excellent quality. However, unsupervised models have been limited by the rigidity of their domain representations. Many methods can only be trained to map between two domains (Zhu et al., 2017; Taigman et al., 2017), or a fixed set of domains (Choi et al., 2018, 2020; Lee et al., 2020). Even models designed to accommodate new domains at test-time either rely on restricting the domain to stylistic features (Liu et al., 2019), or requiring re-training for every new example (Benaim & Wolf, 2018). Moreover, to the best of our knowledge, no model has the capacity to process sets of examples specifying both the source and target domain at test-time, but rely either on explicit conditioning or on one single example. All these constraints limit the model's ability to understand unseen domains and transfer knowledge between them.

Conversely, state-of-the-art methods to perform novel view synthesis rely either on structural assumptions about the geometry of the scene (Sitzmann et al., 2019; Yoon et al., 2020) or on explicit conditioning on camera viewpoint (content) (Eslami et al., 2018; Mildenhall et al., 2020). This restricts the model's usefulness when the viewpoint or scene structure is missing or difficult to describe explicitly.

Perhaps the closest inspiration for our approach comes from the influential work on Semi-Supervised learning by Kingma et al. (2014), where they design a generative model with two latent variables: class (domain) and z (content). They recognize that a limitation of their method is that the number of generative likelihood evaluations scales linearly with the number of classes, since the class variable is categorical.

1.2 Our contribution

We wish to address a general formulation of the domain-content problem that does not rely on any explicit conditioning or constraints on the nature of the disentangled features. We treat domain and content as independent continuous random variables. The strength of our model is that it allows for the specification of new domains at test-time by provision of any number of examples, and permits content queries through examples as well. The continuous nature of the latent representations is also useful for measuring similarity between domains, or for classifying inputs by content irrespective of domain.

In this work:

- We propose a probabilistic model of domain-content with an associated neural architecture built upon the paradigm of the Variational Autoencoder. Our model has the capacity to separate domain and content features in an unsupervised, few-shot manner. The groupwise domain encoder enables it to process unseen domains at test-time, while its novel domain-confusion loss prevents domain and content information from mixing in the latent representations during training.
- We demonstrate qualitatively the ability of our model to perform a generalized task called domain-content fusion, bringing together image-to-image translation and novel view synthe-



Figure 1: Bayesian networks comparing VAE (left) and our Domain-Content model (right).

sis, that requires the model to generate images by combining the domain of one input with the content of another.

• We measure quantitatively the robustness of our disentanglement by ascertaining how well our model's latent representations can predict domain and content features in the data. We record improvements over other disentanglement methods.

2 Probabilistic Domain-Content

We consider a very general formulation of the domain-content problem: Let X be a dataset of images divided into N packs $\underline{\mathbf{x}}_i$ (we use the underline notation to denote a pack of elements), where $i \in \{1 : N\}$. Each pack *i* consists of K_i images $\underline{\mathbf{x}}_i = \{\mathbf{x}_{ik} \mid k \in \{1 : K_i\}\}$. By definition, all elements within a pack belong to the same domain.

Our goal is to create a probabilistic model that exploits the pack structure of the dataset in order to extract a domain random variable \mathbf{m}_i and a pack of content random variables $\underline{\mathbf{o}}_i$ from a given pack of images $\underline{\mathbf{x}}_i$. In order to constitute useful disentangled representations, the inferred latent variables should satisfy certain intuitive principles:

- 1. The representation should be rich enough that one could recover an accurate estimate of the input \underline{x}_i given the distribution of the latent variables \mathbf{m}_i and $\underline{\mathbf{o}}_i$.
- 2. Each image in the pack $\underline{\mathbf{x}}_i$ should have the same associated domain variable \mathbf{m}_i (we have enforced this by construction).
- 3. The distribution of individual content variables o_{ik} should be independent of the domain of the input.

In this work, we propose a probabilistic model with an associated neural architecture that, by following these principles, achieves effective and robust disentanglement of the domain and content factors. Our model follows the Variational Autoencoder paradigm (Kingma & Welling, 2014; Rezende et al., 2014), wherein the latent inference density is used as a sampling distribution for training a generative model. We have been particularly inspired by the semi-supervised approach of Kingma et al. (2014), who have also built a Variational Autoencoder with two latent variables.

2.1 Parametric Generative Model

Our generative model comprises a family of parametric densities p over the variables $\underline{\mathbf{x}}_i$, \mathbf{m}_i and $\underline{\mathbf{o}}_i$. The joint distribution of a pack factorizes according to the Bayesian network in Figure 1:

$$p(\mathbf{x}_i, \mathbf{m}_i, \underline{\mathbf{o}}_i) = p(\mathbf{m}_i) \prod_{k=1}^{K_i} p(\mathbf{o}_{ik}) p(\mathbf{x}_{ik} | \mathbf{m}_i, \mathbf{o}_{ik})$$
(1)

Notice that, when conditioned on \mathbf{m}_i , an individual image variable \mathbf{x}_{ik} is independent of all the other images in the pack $\underline{\mathbf{x}}_i \setminus {\mathbf{x}_{ik}}$ and their corresponding contents $\underline{\mathbf{o}}_i \setminus {\mathbf{o}_{ik}}$. We assign parameters α to the prior over the domain $p_{\alpha}(\mathbf{m}_i)$, β to the prior over the content $p_{\beta}(\mathbf{o}_{ik})$ and θ to the generator density $p_{\theta}(\mathbf{x}_{ik}|\mathbf{m}_i,\mathbf{o}_{ik})$. Only θ is a trainable parameter, since it corresponds to the parameters of our decoder network. Its maximum likelihood estimator is:



Figure 2: Neural architecture of the MO model.

$$\theta' = \arg\max_{\theta} \frac{1}{N} \sum_{i=1}^{N} \log p_{\theta,\alpha,\beta}(\underline{\mathbf{x}}_i), \text{ where}$$
(2)

$$p_{\theta,\alpha,\beta}(\mathbf{\underline{x}}_i) = \mathbb{E}_{p_{\alpha}(\mathbf{m}_i)} \prod_{k=1}^{K_i} \mathbb{E}_{p_{\beta}(\mathbf{o}_{ik})} \log p_{\theta}(\mathbf{x}_{ik} | \mathbf{m}_i, \mathbf{o}_{ik})$$
(3)

2.2 Variational Inference in the Domain-Content Model

Optimizing the likelihood under this formulation would require sampling over the priors $p_{\alpha}(\mathbf{m}_i)$ and $p_{\beta}(\mathbf{o}_{ik})$, a procedure that would converge extremely slowly and leave us with no tractable posterior over the latents conditioned on the images. We, therefore, introduce a parametric inference density $q(\mathbf{m}_i, \mathbf{o}_i | \mathbf{x}_i)$ over which to perform importance sampling during training for faster convergence and a tractable inference posterior. This inference density comprises the domain-content extractor that we set out to create.

According to Kahn & Marshall (1953), the generative latent posterior $p_{\theta,\alpha,\beta}(\mathbf{m}_i, \mathbf{o}_i | \mathbf{x}_i)$ is, itself, the optimal inference density with respect to reducing the variance of the maximum likelihood estimator. We seek, therefore, to design an inference density $q(\mathbf{m}_i, \mathbf{o}_i | \mathbf{x}_i)$ that preserves the conditional relationships between the variables in the generative model, in order to be theoretically capable of recovering the generative posterior. A range of choices are available on how to factorize the inference posterior while still retaining the aforementioned conditional relationships. We choose to condition the content on the domain in order to exploit the conditional independence of contents in a pack given its domain $q(\mathbf{o}_i | \mathbf{m}_i, \mathbf{x}_i) = \prod_{k=1}^{K_i} q(\mathbf{o}_{ik} | \mathbf{m}_i, \mathbf{x}_{ik})$. The inference model, depicted in Figure 1, becomes:

$$q(\mathbf{m}_i, \underline{\mathbf{o}}_i | \underline{\mathbf{x}}_i) = q(\mathbf{m}_i | \underline{\mathbf{x}}_i) \prod_{k=1}^{K_i} q(\mathbf{o}_{ik} | \mathbf{x}_{ik}, \mathbf{m}_i)$$
(4)

We assign parameters ζ to the inference posterior of the domain $q_{\zeta}(\mathbf{m}_i | \mathbf{x}_i)$, and ξ to the posterior inference over the content $q_{\xi}(\mathbf{o}_{ik} | \mathbf{x}_{ik}, \mathbf{m}_i)$. Both ζ and ξ are trainable parameters, since the goal is to learn an inference density accurate enough to recover the image input (Principle 1). They correspond to the parameters of our domain and content encoder.

By applying Importance Sampling, followed by Jensen's Inequality, to the maximum likelihood objective, we obtain the Evidence Lower Bound (ELBO) for our model (the full derivation is available in Appendix A):

Algorithm 1 MorphOus

Input: Dataset **X** formed of N packs each of K_i images, hyperparameter $\lambda = 100$. for i = 1 to N do Randomly select two packs $\underline{\mathbf{x}}_i$ and $\underline{\mathbf{x}}_i$ $\mathbf{m}_i \sim q_{\zeta}(\mathbf{m}_i | \underline{\mathbf{x}}_i)$ $L_i^m := \log q_\zeta(\mathbf{m}_i | \mathbf{x}_i) - \log p_\alpha(\mathbf{m}_i)$ for k = 1 to K_i do $\mathbf{o}_{ik} \sim q_{\xi}(\mathbf{o}_{ik}|\mathbf{x}_{ik},\mathbf{m}_i)$ $L_{ik}^{o} := \log q_{\xi}(\mathbf{o}_{ik}|\mathbf{x}_{ik},\mathbf{m}_{i}) - \log p_{\beta}(\mathbf{o}_{ik})$ $L_{ik}^{r} := (\mathbf{x}_{ik} - G(\mathbf{m}_{i},\mathbf{o}_{ik}))^{2}$ end for $\begin{array}{l} L_{i} = L_{i}^{m} + \sum_{k=1}^{K_{i}} [L_{ik}^{r} + L_{ik}^{o}] \\ \textit{//repeat for pack } \underline{\mathbf{x}}_{j} \end{array}$ $\begin{array}{l} L_{j} = L_{j}^{m} + \sum_{k=1}^{K_{j}} [L_{jk}^{r} + L_{jk}^{o}] \\ \textit{// domain-confusion loss} \end{array}$ $L^c := \text{DomConfLoss}(\mathbf{o}_i, \mathbf{o}_i)$ $L := L_i + L_j + \lambda L^c$ // update parameters $\theta, \zeta, \xi := \operatorname{Adam}(L)$ end for

Algorithm 2 DomConfLoss

Input: Two packs of content variables $\underline{o}_i, \underline{o}_j$ of size K_i and K_j . **Output:** Loss value verifying whether the packs have the same distribution. // choose random sizes for pack splits $A_i \sim \mathcal{U}(1, K_i), B_i := K_i - A_i$ $A_j \sim \mathcal{U}(1, K_j), B_j := K_j - A_j$ $\underline{a}_i, \underline{b}_i := \text{split}(\underline{o}_i, A_i, B_i)$ $\underline{a}_j, \underline{b}_j := \text{split}(\underline{o}_j, A_j, B_j)$ // get "real" discriminator predictions $\mathbf{r}_i := D(\sum_{k=1}^{A_i} H(\mathbf{a}_{ik}), \sum_{l=1}^{B_i} H(\mathbf{b}_{ll}))$ $r_j := D(\sum_{k=1}^{A_i} H(\mathbf{a}_{ik}), \sum_{l=1}^{B_j} H(\mathbf{b}_{jl}))$ // get "fake" discriminator predictions $\mathbf{f}_a := D(\sum_{k=1}^{A_i} H(\mathbf{a}_{ik}), \sum_{l=1}^{B_j} H(\mathbf{b}_{jl}))$ $f_b := D(\sum_{k=1}^{B_i} H(\mathbf{b}_{ik}), \sum_{l=1}^{B_j} H(\mathbf{b}_{jl}))$ $L^c := \log \mathbf{r}_i \mathbf{r}_j + \log(1 - \mathbf{f}_a)(1 - \mathbf{f}_b)$ // update discriminator parameters $D, H := \text{Adam}(-L^c)$ **return** L^c

$$\log p_{\theta,\alpha,\beta}(\underline{\mathbf{x}}_i) \ge \mathbb{E}_{q_{\zeta}(\mathbf{m}_i|\underline{\mathbf{x}}_i)} \sum_{k=1}^{K_i} \mathbb{E}_{q_{\xi}(\mathbf{o}_{ik}|\mathbf{x}_{ik},\mathbf{m}_i)} \log p_{\theta}(\mathbf{x}_{ik}|\mathbf{m}_i,\mathbf{o}_{ik}) \quad (\text{reconstruction}) \quad (5)$$

$$-D_{\mathrm{KL}}[q_{\zeta}(\mathbf{m}_{i}|\mathbf{\underline{x}}_{i}) || p_{\alpha}(\mathbf{m}_{i})]$$
 (domain) (6)

$$-\mathbb{E}_{q_{\zeta}(\mathbf{m}_{i}|\mathbf{x}_{i})}\sum_{k=1}^{K_{i}}D_{\mathrm{KL}}[q_{\xi}(\mathbf{o}_{ik}|\mathbf{x}_{ik},\mathbf{m}_{i}) \mid \mid p_{\beta}(\mathbf{o}_{ik})] \quad \text{(content)}$$
(7)

This separates neatly into a reconstruction loss and two regularization penalties, for the domain and content variables. The expression of the reconstruction loss optimizes Principle 1 directly, as it encourages precise estimates of the output image.

2.3 Neural Architecture

Following the VAE paradigm, we implement the three trainable parametric densities of our model as three normal distributions with diagonal covariance, whose mean and variance are computed by feed-forward neural architectures. The generator density takes the form of a normal distribution with fixed variance and mean computed by the generator network G, taking as input the concatenated domain and content codes. In practice, the output image will be the mean of the distribution, rather than samples from it.

$$p_{\theta}(\mathbf{x}_{ik}|\mathbf{m}_i, \mathbf{o}_{ik}) = \mathcal{N}(\mu_x, 1), \text{ where } \mu_x = G(\mathbf{m}_i, \mathbf{o}_{ik})$$
 (8)

The domain inference density is a normal that requires its parameters to be computed by a neural architecture processing a variable number of un-ordered, exchangeable inputs. For this, we use a Deep Set network architecture (Zaheer et al., 2017), whereby each input is individually encoded by a the same network E_m , then the outputs are averaged together, and the result is passed through a second network F_m . We average the outputs instead of summing them (as used in Zaheer et al.)



Figure 3: Neural architecture of the Domain-Confusion discriminator η .

(2017)) because we want the inference density of the domain to be agnostic to the number of inputs in the pack.

$$q_{\zeta}(\mathbf{m}_i|\underline{\mathbf{x}}_i) = \mathcal{N}(\mu_m, \sigma_m^2), \text{ where } \mu_m, \sigma_m = F_m\left(\frac{1}{K_i}\sum_{k=1}^{K_i} E_m(\mathbf{x}_{ik})\right)$$
(9)

The content inference density is a normal with parameters computed by encoding an image with a network E_o , then concatenating the output with the domain code of the pack, then passing it through another network F_o .

$$q_{\xi}(\mathbf{o}_{ik}|\mathbf{x}_{ik},\mathbf{m}_{i}) = \mathcal{N}(\mu_{o},\sigma_{o}^{2}), \text{ where } \mu_{o},\sigma_{o} = F_{o}(\mathbf{m}_{i},E_{o}(\mathbf{x}_{ik}))$$
(10)

Diagrams depicting each of these architectures are displayed in Figure 2. We employ the reparametrization trick (Kingma & Welling, 2014) to sample from the inference posterior over the latents. As for the domain and content priors, they are fixed arbitrary normal distributions with mean 0 and variance 1. A more complete specifications of the neural implementation is available in Appendix B.

3 Domain-Confusion Loss

So far, Principle 1 is encouraged by optimizing the Evidence Lower Bound, while Princicple 2 is enforced by construction. We inspect whether the model also satisfies Principle 3, which requires that the distribution of inferred content variables in a pack be orthogonal to the domain of the pack. We can reformulate this principle more precisely to claim that, in the limit of infinitely large packs, the distribution of a random variable \tilde{o} , denoting the random choice of one content variable o_{ik} from a pack of inferred content variables o_i , should be the same same regardless of the pack of origin.

$$q_{\zeta,\xi}(\tilde{\mathbf{o}}|\mathbf{x}_i) = q_{\zeta,\xi}(\tilde{\mathbf{o}}|\mathbf{x}_i), \text{ where } K_i, K_i \to \infty$$
(11)

This statement relies on the fact that, when the size of the pack tends to infinity, the random picking of one inferred content variable from the pack defines a distribution over content values conditioned on the "true" domain of the pack. Since the probability of content features should be independent of domain features, the empirical density $q(\tilde{\mathbf{o}}|\mathbf{x}_i)$ should also stay fixed as *i* changes.

Claim When the Domain-Content Evidence Lower Bound (equation 5) is maximized to its theoretical potential, then equation 11 is satisfied. In other words, when the inference latent posterior $q_{\zeta,\xi}(\mathbf{m}_i, \mathbf{o}_i | \mathbf{x}_i)$ approaches the generative latent posterior $p_{\theta,\alpha,\beta}(\mathbf{m}_i, \mathbf{o}_i | \mathbf{x}_i)$, and the generative data likelihood $p_{\theta,\alpha,\beta}(\mathbf{x}_i)$ produces samples indistinguishable from the real data, then the distribution of content variables will become independent of the "true" domain of the pack, in the limit of large packs. A discussion and proof of this is included in Appendix A.

This result shows that our probabilistic model is, in theory, sufficient to satisfy Principle 3. However, this state cannot be achieved in practice, because of architectural limitations on both the generative and inference density families. Therefore, in order to encourage the realization of Principle 3, we can constrain the space of inference densities to those for which equation 11 holds, at least approximately. This need is reinforced by empirical observations of the unconstrained model, which reveal that the distribution of inferred content variables within a pack is highly sensitive to the task, network architecture and choice of hyperparameters.

In this work, we tackle this problem practically by proposing an adversarial loss which encourages the homogeneous distribution of content variables across packs by penalizing differences between pairs of packs in their set of content values. We call this the Domain-Confusion loss, and we show empirically in Table 1 that it increases the quality and robustness of the disentanglement.

The loss is built around an adversarial discriminator η which receives as input a pair of packs of content values and is trained to output 1 if the two packs come from the same distribution, or 0 if they are differently distributed. This is called a verification task, inspired by Sohn et al. (2019), and we apply it to contrast pairs of sub-packs coming from the same pack with pairs of sub-packs coming from different packs. Concretely, every iteration of training takes as input two packs of content values $\underline{o}_i, \underline{o}_j$, which we split randomly into $(\underline{a}_i, \underline{b}_i)$ and $(\underline{a}_i, \underline{b}_i)$, respectively. The loss takes the value:

$$L^{c}(\underline{\mathbf{o}}_{i},\underline{\mathbf{o}}_{j}) = \log \eta(\underline{\mathbf{a}}_{i},\underline{\mathbf{b}}_{i}) + \log \eta(\underline{\mathbf{a}}_{j},\underline{\mathbf{b}}_{j}) + \log(1 - \eta(\underline{\mathbf{a}}_{i},\underline{\mathbf{a}}_{j})) + \log(1 - \eta(\underline{\mathbf{b}}_{i},\underline{\mathbf{b}}_{j}))$$
(12)

The greater the loss, the more the discriminator can distinguish between the two packs. Because the architecture of η needs to accommodate two packs of varying sizes, we implement it as a Deep Set (Zaheer et al., 2017), just like in the case of the domain encoder. Unlike to the domain encoder, here we use summation instead of averaging, since η does not need to be agnostic to the number of inputs. The discriminator takes the form:

$$\eta(\underline{\mathbf{o}}_i, \underline{\mathbf{o}}_j) = D\left(\sum_{k=1}^{K_i} H(\mathbf{o}_{ik}), \sum_{l=1}^{K_j} H(\mathbf{o}_{jl})\right)$$
(13)

where D and H are neural networks. A diagram of this architecture is displayed in Figure 3. The full implementation is available in Algorithm 2.

4 **Experiments**

We evaluate our model on the generalized domain-content fusion task mentioned in the introduction. Given a trained model and two unseen image packs $\underline{\mathbf{x}}_i$, $\underline{\mathbf{x}}_j$, we extract the domain of $\underline{\mathbf{x}}_i$ and the contents of every image in $\underline{\mathbf{x}}_j$, and then use them to generate K_j output images.

This task enables us to visually inspect the quality of the disentanglement by judging how well the model follows each of the domain-content principles: Are the images of high-quality? Do they have the same domain features? Do the content features of each correspond to those of the associated input image?

We apply our model to three datasets: a dataset of font images collected from Google Fonts, the Small Norb Dataset (LeCun et al., 2004) and an original dataset, called Silhouettes, comprising 3-dimensional block shapes imaged at various rotation angles. Detailed descriptions of the datasets are available in Appendix C. For the Silhouettes and Google Font datasets, we test only on domains which have been withheld during training, in order to show how well the model generalizes to new domains. We display results for 5 testing packs of each dataset in Figure 4. Further results on more datasets are available in the Appendix D. We also provide results in Appendix D for testing the model on both unseen domains and unseen contents. The results appear to separate domain and content features very well without sacrificing the image quality as compared to the VAE.

4.1 Predicting ground-truth factors from the latent representation

In order to obtain more quantitative evidence of disentanglement, we adapt the factor regression method also used by Greff et al. (2019) to measure factor information and interpretability in the



(c) Disentangling object shape and object rotation.

Figure 4: Qualitative results of our model on the fusion task are displayed with a black border. Each row represents a domain and each column represents a content. The topmost row is the content input and each row of the leftmost 5 columns is a domain input.

Table 1: How we	ll can the latent repre	esentation predict the	ground-truth factors
14010 11 11011 110	i can nic iacont iopie	beindarion predict die	

	MO (w/ DC)		MO (w/o DC)				
Factor (metric)	domain	content	domain	content	FactorVAE	VAE	Guessing
Object shape (CE) Rotation (MSE)	-0.023 673.4	-0.449 456.2	-0.051 672.8	-0.402 533.2	-0.211 563.1	-0.236 597.8	-0.451 671.3

latent space. The method involves learning a simple linear mapping between the latent space of the trained model and the value space of the ground-truth factor. A high predictive accuracy implies that the latent representation contains the necessary factor information and also organizes it in an easily interpretable way.

We complete this experiment on our Silhouettes dataset, and learn separate predictors for the domain features (object shape) and content features (rotation angle). Predicting the shape is a 27-way binary classification and predicting the rotation angle is a two-way regression. We provide more details on the encoding of these features in Appendix C. We compare each latent representation of our model (domain and content) with a VAE, a FactorVAE (Kim & Mnih, 2018), considered to be state-of-the-art in disentanglement, and random guessing. In the case of our model, the goal is that each of the two representations should predict its own factor as much as possible, and not to predict the other's factor. We include a comparison of our model with and without the Domain-Confusion loss as an ablation study on the impact of this loss. Details on the experimental setup and measurements, as well as comparisons with other models on other datasets, are available in Appendix D.

The results, displayed in Table 1, reveal not only that the latent variables of our model predict their corresponding factors far better than the FactorVAE or VAE, but also that the cross-over predictions

are no better than random. This result shows that domain and content features are concentrated successfully in their designated representations. Moreover, we can see a marked improvement in the model with the Domain-Confusion loss over the one without it, in the case of both domain and content factors.

5 Conclusion

In this work, we have described a general problem of disentangling the domain and content generative factors, and proposed a probabilistic model with an associated neural network to solve this task. We have built the model according to the VAE paradigm and introduced the Domain-Confusion loss to compensate for limitations brought upon by the neural architecture. We have proven the effectiveness of our disentanglement solution by providing both qualitative results on the domain-content fusion tasks, and quantitative measures of the predictive power of the latent representations.

One crucial direction to explore in future work is the relationship between domain-content disentanglement and the notion of invariant risk Arjovsky et al. (2019) in the context of causal inference. Designing methods that can separate confounding environmental factors (domain) from the factors of interest (content) would lead to significant innovations in many research fields, especially in the study of medical counterfactuals. The method presented here is not yet able to perform such separation in the case where content distributions vary across domains, since it uses the homogeneity across domains as a proxy for identifying content factors.

6 Broader Impact

One of the main ethical faults associated with the application of statistical learning to real-world problems is the acquisition of any biases that might be present in the training dataset. As Arjovsky et al. (2019) discuss in their work, classical deep learning methods minimizing the expected risk of their hypothesis cannot distinguish between spurious correlations and true mechanisms. Therefore, naive statistical correspondences are drawn between phenomena that are not causally connected, creating unintended consequences with potential scientific or social impact. Although our model is still an expected-risk minimizing algorithm, we note that it could also be used as a paradigm for diagnosing dataset biases. For example, if two distinct populations of elements, sharing the same set of classes in a classification task, vary in their representation of different classes across different datasets, then the same individual placed in various datasets will produce, in turn, different content encodings, revealing the biases in the dataset. Our model is still very theoretical, but we believe it is a step towards a deeper study of the relationship between element and environment in the context of deep learning.

References

- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- Benaim, S. and Wolf, L. One-shot unsupervised cross domain translation. In Advances in Neural Information Processing Systems, pp. 2104–2114, 2018.
- 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.
- Burgess, C. P., Matthey, L., Watters, N., Kabra, R., Higgins, I., Botvinick, M., and Lerchner, A. Monet: Unsupervised scene decomposition and representation. arXiv preprint arXiv:1901.11390, 2019.
- Choi, Y., Choi, M., Kim, M., Ha, J.-W., Kim, S., and Choo, J. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8789–8797, 2018.

- Choi, Y., Ha, Y., Yoo, J., Uh, J.-W., Kim, S., and Choo, J. Stargan v2: Diverse image synthesis for multiple domains. In *Computer Vision and Pattern Recognition (CVPR)*, 2020.
- Engelcke, M., Kosiorek, A. R., Jones, O. P., and Posner, I. Genesis: Generative scene inference and sampling with object-centric latent representations. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=BkxfaTVFwH.
- Eslami, S. A., Rezende, D. J., Besse, F., Viola, F., Morcos, A. S., Garnelo, M., Ruderman, A., Rusu, A. A., Danihelka, I., Gregor, K., et al. Neural scene representation and rendering. *Science*, 360 (6394):1204–1210, 2018.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Gonzalez-Garcia, A., van de Weijer, J., and Bengio, Y. Image-to-image translation for cross-domain disentanglement. In Advances in Neural Information Processing Systems, pp. 1287–1298, 2018.
- Greff, K., Kaufman, R. L., Kabra, R., Watters, N., Burgess, C., Zoran, D., Matthey, L., Botvinick, M., and Lerchner, A. Multi-object representation learning with iterative variational inference. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pp. 2424–2433. PMLR, 2019.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.
- Kahn, H. and Marshall, A. W. Methods of reducing sample size in monte carlo computations. *Journal* of the Operations Research Society of America, 1(5):263–278, 1953.
- Kim, H. and Mnih, A. Disentangling by factorising. In *Proceedings of the 35rd International Conference on Machine Learning*, volume 80, pp. 2649–2658, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. 2nd International Conference on Learning Representations, ICLR 2019, 2014.
- Kingma, D. P., Mohamed, S., Rezende, D. J., and Welling, M. Semi-supervised learning with deep generative models. In Advances in neural information processing systems, pp. 3581–3589, 2014.
- LeCun, Y., Huang, F. J., and Bottou, L. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2004. CVPR 2004., volume 2, pp. II–104. IEEE, 2004.
- Lee, H.-Y., Tseng, H.-Y., Mao, Q., Huang, J.-B., Lu, Y.-D., Singh, M., and Yang, M.-H. Drit++: Diverse image-to-image translation via disentangled representations. *International Journal of Computer Vision*, 2020.
- Liu, M.-Y., Huang, X., Mallya, A., Karras, T., Aila, T., Lehtinen, J., and Kautz, J. Few-shot unsupervised image-to-image translation. In *IEEE International Conference on Computer Vision* (*ICCV*), 2019.
- Locatello, F., Bauer, S., Lucic, M., Rätsch, G., Gelly, S., Schölkopf, B., and Bachem, O. Challenging common assumptions in the unsupervised learning of disentangled representations. *International Conference on Machine Learning*, 2019.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. Nerf: Representing scenes as neural radiance fields for view synthesis, 2020.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 1278–1286. PMLR, 2014.
- Sitzmann, V., Zollhöfer, M., and Wetzstein, G. Scene representation networks: Continuous 3dstructure-aware neural scene representations. In Advances in Neural Information Processing Systems, pp. 1119–1130, 2019.

- Sohn, K., Shang, W., Yu, X., and Chandraker, M. Unsupervised domain adaptation for distance metric learning. In *International Conference on Learning Representations*, 2019. URL https: //openreview.net/forum?id=BklhAj09K7.
- Taigman, Y., Polyak, A., and Wolf, L. Unsupervised cross-domain image generation. In 5th International Conference on Learning Representations, ICLR, 2017.
- Tschannen, M., Bachem, O., and Lucic, M. Recent advances in autoencoder-based representation learning. In *Third workshop on Bayesian Deep Learning (NeurIPS 2018)*, 2018. URL http: //www.nari.ee.ethz.ch/pubs/p/autoenc2018.
- van Steenkiste, S., Locatello, F., Schmidhuber, J., and Bachem, O. Are disentangled representations helpful for abstract visual reasoning? In *Advances in Neural Information Processing Systems*, pp. 14222–14235, 2019.
- Yoon, J. S., Kim, K., Gallo, O., Park, H. S., and Kautz, J. Novel view synthesis of dynamic scenes with globally coherent depths from a monocular camera. In *IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2020.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *Advances in neural information processing systems*, pp. 3391–3401, 2017.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Upaired image-to-image translation using cycleconsistent adversarial networks. In *Proceedings of the IEEE international conference on computer* vision, pp. 2223–2232, 2017.