

Disentangling Generative Factors in Natural Language with Discrete Variational Autoencoders

Giangiaco^o Mercatali^{*}
University of Manchester

Andr  Freitas[†]
Idiap Research Institute
University of Manchester

Abstract

The ability of learning disentangled representations represents a major step for interpretable NLP systems as it allows latent linguistic features to be controlled. Most approaches to disentanglement rely on continuous variables, both for images and text. We argue that despite being suitable for image datasets, continuous variables may not be ideal to model features of textual data, due to the fact that most generative factors in text are discrete. We propose a Variational Autoencoder based method which models language features as discrete variables and encourages independence between variables for learning disentangled representations. The proposed model outperforms continuous and discrete baselines on several qualitative and quantitative benchmarks for disentanglement as well as on a text style transfer downstream application.

1 Introduction

A fundamental challenge in Natural Language Processing (NLP) is being able to control generative factors for text, such as tense, gender, negation, which are characterised by an entangled representation in traditional neural networks, making it difficult to control them. Disentangled representation learning aims to provide an interpretable representation of latent features, and a framework for controlling the change of specific features, by separating distinct generative factors in the data (Bengio et al., 2013).

Various disentanglement approaches for neural networks have been proposed. Chen et al. (2016) achieve disentanglement with a Generative Adversarial Network (GAN), by maximizing the mutual information between latent variables and generated samples, while (Higgins et al., 2016; Burgess et al., 2018; Kim and Mnih, 2018) fine-tune the param-

^{*}name.surname@postgrad.manchester.ac.uk

[†]name.surname@manchester.ac.uk

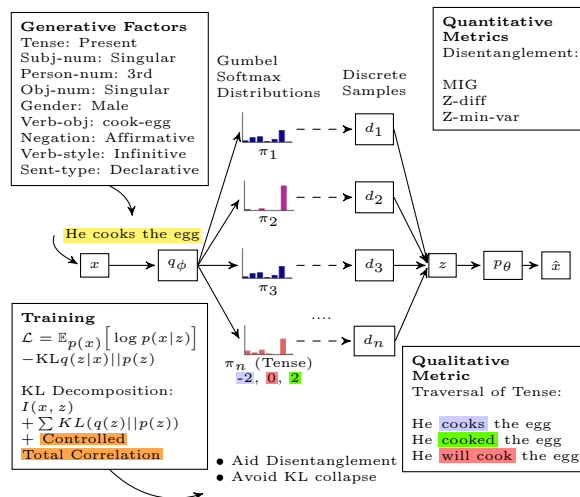


FIGURE 1: Overview of the proposed Discrete Controlled Total Correlation (DCTC) model: The KL decomposition encourages independence between variables, which are encoded as discrete latents, to capture the dimensions of linguistic features. The representation is probed with latent traversals and quantitative metrics.

eter which controls the KL divergence in a Variational AutoEncoder (VAE) (Kingma and Welling, 2014). In the NLP domain, John et al. (2019) use adversarial losses to separate the style and content embeddings, while Cheng et al. (2020) disentangle style and content embeddings by minimizing their mutual information.

After a thorough review of the literature, we find that there are currently two main issues in the area of disentanglement in NLP. First, the mentioned approaches operate using Gaussian distributions in a continuous space. Although a continuous representation may be suitable to encode images, for text data one should rather specifically consider discrete distributions of the feature set. In fact, generative features of a sentence mostly belong to a discrete domain, for example, one would encode the gender feature either as male or female, while the tense can be modelled as having three values,

such as present, past, future. Secondly, while in the realm of Representation Learning, numerous quantitative and qualitative evaluation methods for disentanglement have been designed, there is still a gap in the adaptation and adoption of these methods into the language domain. As a result, despite being a critical step for enabling the interpretability of NLP models, text disentanglement is still not fully explored.

We address the highlighted problems as follows:

1) We design the first VAE-based architecture (shown in Figure 1), where linguistic features are encoded as discrete latent variables via the Gumbel-Softmax trick (Jang et al., 2016) and disentanglement is jointly enforced in the objective function. We derive a decomposition of the VAE evidence lower bound (ELBO), where independence between variables is encouraged by fine-tuning the Total Correlation term, with the goal of achieving disentanglement. At the same time, an information channel mechanism is introduced, and increased incrementally during training, to avoid the issue of the posterior collapse (Bowman et al., 2016).

2) We provide, to our knowledge, the first extensive evaluation of the text representation, under the lenses of contemporary disentanglement methods. We propose to probe the quality of the representation by traversing and decoding each latent, expecting a disentangled representation to show only a single dimensional change (for example in tense, as in Figure 1). In addition, we show that the proposed model outperforms numerous baselines on three quantitative disentanglement metrics from the Representation Learning literature. Finally, we show the beneficial effect of the proposed representations in the task of text style transfer.

2 Background

Language VAEs. Language VAEs (Bowman et al., 2016) are generative models widely used in NLP tasks such as text style transfer (John et al., 2019) and conditional text generation (Cheng et al., 2020). The VAE architecture consists of a decoder network p_θ and encoder network q_ϕ , acting over a sequence of sentences x_1, \dots, x_n .

The VAE considers a multivariate Gaussian prior distribution $p(z)$ and generate a sentence x with the decoder $p_\theta(x|z)$. The joint distribution for the decoder is defined as $p(z)p_\theta(x|z)$, which, for a sequence of tokens x of length T result as $p_\theta(x|z) = \prod_{i=1}^T p_\theta(x_i|x_{<i}, z)$.

The framework’s objective is to maximize the expectation over the full dataset of the log-likelihood, in other words, $\mathbb{E}_{p(x)} \log p_\theta(x)$. However, this term is intractable, and thus the variational distribution q_θ is introduced to approximate $p_\theta(z|x)$. As a result, an evidence lower bound \mathcal{L}_{VAE} (ELBO) where $\mathbb{E}_{p(x)} [\log p_\theta(x)] \geq \mathcal{L}_{\text{VAE}}$, is derived as follows:

$$\mathcal{L}_{\text{VAE}} = \underbrace{\mathbb{E}_{q_\phi(z|x)} \left[\log p_\theta(x|z) \right]}_{\textcircled{i}} - \underbrace{\text{KL}q_\phi(z|x) || p(z)}_{\textcircled{ii}} \quad (1)$$

In Eq. (1), \textcircled{i} represents the reconstruction error, which encourages the model to encode the data x into the latent variables z , while \textcircled{ii} represents the KL divergence regularization term, which forces the variational distribution $q_\phi(z|x)$ to be similar to the prior $p(z)$.

Disentanglement in VAEs. The key intuition of a *disentangled* representation is that it should separate the independent factors of variation in the data. As a result, a change in a single underlying factor is reflected into a change in a single factor of the learned representation, while being invariant to changes in other factors (Bengio et al., 2013).

Early approaches to disentanglement are GAN-based (Chen et al., 2016), however, most contemporary methods focus on the VAE framework, which improves the training stability. The β VAE framework (Higgins et al., 2016) introduces a hyperparameter β for tuning the KL term \textcircled{ii} in the ELBO (1) (i.e. $\mathcal{L}_{\beta\text{VAE}} = \textcircled{i} - \beta\textcircled{ii}$) and demonstrates that on one hand, a $\beta > 1$ leads the model to learn disentangled representations, however, the disentanglement comes at the cost of a low reconstruction fidelity.

Recent approaches (Kim and Mnih, 2018; Chen et al., 2018) demonstrate that disentanglement can be achieved without sacrificing the reconstruction fidelity, by decomposing the KL term in the ELBO, and tuning only the terms that encourages independence between latents, while not penalizing the mutual information between latents and data.

3 Proposed Approach

Model Overview. Differently from previous approaches to disentanglement (Higgins et al., 2016; Kim and Mnih, 2018; Chen et al., 2018), we focus

our efforts into leveraging the discrete generative factors present in natural language, and design a framework, which we name Discrete Controlled Total Correlation (DCTC), where language factors are encoded as discrete latent variables, while the representation is enforced to be disentangled. We first reformulate the continuous VAE for discrete variables by providing a suitable discrete reparameterization process. We then decompose the ELBO, and provide a modification of the Total Correlation term for controlling disentanglement.

3.1 Linguistic Features as Discrete Variables

Discrete Variables. Discrete representations present strategic advantages for modeling natural language. Firstly, a discrete encoding can capture the dimensions of a linguistic feature (e.g. tense is a three values variable). In addition, since the dimensions of language features are known, it is possible to induce semantic factor biases from them, and thus enhance disentanglement (Locatello et al., 2019).

In order to model the discrete generative factors of a sentence or passage, we first define a set of discrete latent variables d , that are sampled from a Gumbel-Softmax distribution (Jang et al., 2016) and a posterior $q_\phi(d|x)$. This is the first step to encode the discrete factors of variation in natural language, such as tense and gender. The continuous VAE in (1) becomes:

$$\mathcal{L} = \mathbb{E}_{q_\phi(d|x)} \left[\log p_\theta(x|d) \right] - \text{KL}(q_\phi(d|x) \parallel p(d)) \quad (2)$$

Discrete Reparameterization. In order to back-propagate with discrete variables, we need to extend the reparameterization-trick from (Kingma and Welling, 2014). In fact, the vanilla VAE only considers a continuous variable c , and $q_\phi(c|x)$ is parameterized by $q_\phi(c|x) = \prod_i q_\phi(c_i|x)$ where each distribution is Gaussian. Since in our loss (2) we consider d , that is a discrete variable, $q_\phi(d|x)$ is non-differentiable, thus we resort to the Gumbel-Softmax trick (Jang et al., 2016), which provides a tool for sampling from a continuous approximation of a discrete distribution. The Gumbel-Softmax trick considers a discrete variable with class probabilities π_1, \dots, π_k , and draws samples g_1, \dots, g_k from a Gumbel distribution, as follows:

$$y_i = \frac{\exp((\log(\pi_k) + g_k)/\tau)}{\sum_{j=1}^k \exp((\log(\pi_j) + g_j)/\tau)} \quad (3)$$

By plugging in the samples g_i and the class probabilities π in Eq. (3), we generate a k -dimensional vector y , that is the continuous approximation of the one-hot-encoded representation of the discrete variable d . In fact, as τ approaches 0, samples from the Gumbel-Softmax distribution become one-hot, making it discrete.

With this approximation mechanism in place, we can define the prior $p(d)$ in Eq. (2) as a product of Gumbel-Softmax distributions, which makes the decoder $q_\phi(d|x)$ differentiable and enables us to train the discrete VAE. For our purpose, each distribution represents a linguistic feature, and can be set to a discrete dimension.

Model Architecture. The model architecture is depicted in Figure 1. We consider a training sample such as a sentence x of length T composed by x_1, \dots, x_T . Our model is built using a LSTM encoder which receives a sentence x , and samples n discrete latent variables; and a LSTM decoder which receives n discrete variables and merges them in a sentence. Each discrete latent variable π_i is sampled from the Gumbel-Softmax distribution, then the Gumbel-Softmax trick (Jang et al., 2016) is used for drawing a discrete sample d_i . Finally, the samples d_1, \dots, d_n are fed to the decoder after being concatenated. The proposed model also aims to learn disentangled representations of sentences. This is achieved with a modification of the ELBO, where we encourage independence between latent variables, as explained in the following section.

3.2 Controlled Total Correlation

ELBO Decomposition Design. For creating our objective function, we consider the KL term in Eq. (2) in expectation over the data, and decompose it, guided by the following considerations. On one hand, we aim to include a Total Correlation (TC) term in our decomposition, because TC is a measure of dependency between variables, and thus, a penalty on it may force the model to find independent factors in the data and strongly aid disentanglement (Kim and Mnih, 2018; Chen et al., 2018). On the other hand, a penalty on a TC term may lead to a KL vanishing issue (Bowman et al., 2016), because it may cause the decoder to ignore the information stored in the latent encoding, leading to poor reconstruction fidelity. Thus, a mechanism to avoid this issue should be integrated in the decomposition.

As a result, we decompose the KL to obtain a

$$\mathbb{E}_{p(x)} \left[\text{KL}q_\phi(d|x) || p(d) \right] = \underbrace{\text{KL}q(d, x) || q(d)p(x)}_{\textcircled{1}} + \underbrace{\sum_j \text{KL}q(d_j) || p(d_j)}_{\textcircled{2}} + \underbrace{\gamma \left| \text{KL}q(d) || \prod_j q(d_j) - C_d \right|}_{\textcircled{3}}$$

FIGURE 2: Discrete Controlled Total Correlation decomposition. The term $\textcircled{3}$ encourages disentanglement.

TC term, for aiding independence, and enhance it by introducing two parameters. First, we add a discrete information capacity C_d , which controls the amount of information that is passed into the TC term. The idea is that by increasing the information channel gradually, starting from zero, we tackle both the KL collapse issue and the high reconstruction loss of the β VAE. Furthermore, we introduce a variable γ , to enforce the TC term to match the discrete information capacity C_d . The final decomposition is shown in Fig. 2, where d_j is the j -th dimension of the latent variable d .

Analysis of Components. The first term $\textcircled{1}$ in Fig. 2 is known as the index-code Mutual Information (Hoffman and Johnson, 2016), and it represents the mutual information between the data and the latent variable. This term has been studied in various disentanglement models, for example (Chen et al., 2016) claim that high mutual information is beneficial for disentanglement, and similarly, Zhao et al. (2019) propose the dropping of the penalty on this term to aid disentanglement. However, this notion is not universally accepted, as (Burgess et al., 2018) showed that a penalty on $\textcircled{1}$ can also encourage disentanglement.

The second term $\textcircled{2}$ prevents latent dimensions from deviating from their priors. However, this term does not have theoretical properties that suggest its utility in the enhancement of disentanglement. For the above reasons we don’t fine-tune the first two terms of the decomposition and focus on the third, which we find to be the most influential for enforcing disentanglement in terms of information theoretic properties.

We name the third term $\textcircled{3}$ the Controlled Total Correlation (CTC). In its original form, the Total Correlation (Watanabe, 1960), is given as $\text{KL}(q(z) || \bar{q}(z))$, where $\bar{q}(z) = \prod_j q(z_j)$, and it represents a generalization of the mutual information, which measures the dependence between variables. Our variation of the TC allows us to control the amount of information encoded in the discrete channel, thus avoiding the collapse of the term.

Training Procedure. To obtain the final loss function of the proposed Discrete Controlled Total Correlation model (DCTC), we consider the KL decomposition described in Fig. 2 and replace it in the original KL term in the discrete VAE ELBO in Eq. (2). The final loss function for our DCTC model is reported in Eq. (4):

$$\mathcal{L} = \mathbb{E}_{q_\phi(d|x)} \left[\log p_\theta(x|d) \right] - \textcircled{1} - \textcircled{2} - \textcircled{3} \quad (4)$$

During training, we increase gradually the discrete information capacity C_d while keeping γ fixed. On one side, this linear increase tackles the known issue of KL vanishing for text VAEs (Bowman et al., 2016), on the other side, it let the model maximise disentanglement. More specifically, for low values of C_d the TC term is collapsed to zero, and the reconstructed sentences are not faithful. However, as C_d is increased, the TC terms start to become greater than zero, and the reconstruction becomes more accurate. This improvement continues until the TC factors for each discrete latent are non-zero, and the reconstruction is identical.

Parameter choice. The choice of the parameters γ and C_d is derived from experimental results and it is guided by some necessary constraints. First, there should not be a tradeoff between a smaller reconstruction error and the information capacity constraint, thus the value of γ needs to be large enough to maintain the capacity at the desired value. Second, after constraining the capacity of the discrete information channel, C_d can be chosen to maximise the capacity of the channel, as a result, the model is prompted to use all the latent variables of the discrete distribution.

4 Related Work

Disentanglement in NLP. We identify two types of approaches. 1) Multiple-losses: Hu et al. (2017); John et al. (2019) encourage disentanglement with adversarial losses for style transfer, while Sha and Lukasiewicz (2021) propose to improve the training stability, using multiple non-adversarial losses. 2) Information-theoretic: Cheng et al. (2020) propose to disentangle style and content by minimizing the mutual information between the latent and

the observed variable, while Colombo et al. (2021) introduce an upper bound of mutual information, showing its benefits in fair classification. Differently from all these approaches, we model linguistic features as discrete variables, which allows us to enforce control on the encodings’ dimensions.

Discrete Latent Variables. Various approaches present discrete encodings for language generation. Shu et al. (2020) enhance control and diversity in generation with latent spaces, while Guo et al. (2020) leverage textual evidence to guide the generation. Both methods are based on the VQVAE (van den Oord et al., 2017), which we consider in our experiments. Bao et al. (2020) encodes discrete latent variables into Transformer blocks for dialogue generation. Differently from these methods, we leverage discrete variables to optimize disentanglement.

5 Experiments

In this section, we evaluate the disentanglement of the proposed model with qualitative and quantitative methods, against several baselines. Furthermore, the benefits of our model’s encodings are demonstrated in the downstream task of text style transfer.

5.1 Qualitative Evaluation

Latent Traversals. After training our model, we can evaluate the disentanglement quality of the representations by analysing the traversals of the latent space. Traversal evaluation is a standard procedure with image disentanglement (Higgins et al., 2016; Kim and Mnih, 2018), but represents a novelty for text datasets.

A visual explanation of how a traversal for textual data works is provided in Fig. 3. The traversal of a latent factor is given by decoding the vectors corresponding to the latent variables, where the evaluated factor is changed within a fixed interval (e.g. [-2, 2]), while all others are kept fixed. If the representation is disentangled, when a latent factor is traversed, the decoded sentences should only change with respect to that factor. This means that after training the model we are able to probe the representation for each latent variable.

Experimental Setup. We evaluate the latent traversals on the dSentences dataset (M’Charrak, 2018), which is composed by 37,000 sentences, and provides the annotations for 9 generative fac-

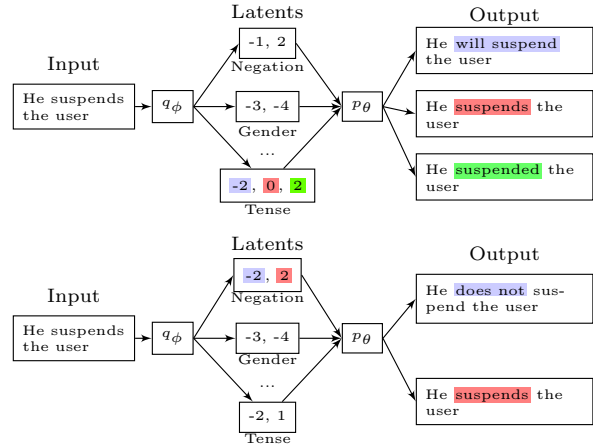


FIGURE 3: Traversals of generative factors. Top: tense. Bottom: negation.

Factor	Dimensions	Values
Verb/object	1100	[Verb/obj variations]
Gender	2	[Male, Female]
Negation	2	[Affirmative, Negative]
Tense	3	[Present, Future, Past]
Subject number	2	[Singular, plural]
Object number	2	[Singular, plural]
Sentence Type	2	[Interrogative, Declarative]
Person number	3	[1st, 2nd, 3rd person]
Verb style	2	[Gerund, Infinitive]

TABLE 1: Generative factors in dSentences data.

tors. Since we know in advance the number and dimensions of the dataset’s generative factors, as reported in Table 1, we set our DCTC model to consider that specific setting, namely 9 discrete generative factors with their true dimensions. The parameter γ from Fig. 2 is fixed at 50 while the discrete channel C_{z_d} is increased from 0 to 30 in 25k steps. Similarly, we set the continuous models to learn 9 Gaussian latent variables, however, differently from the discrete case, it is not possible to map the continuous latent to a dimension.

Baselines. We compare our model against several types of state-of-the-art VAEs. 1) Continuous disentanglement models, such as β VAE (Higgins et al., 2016), CCI-VAE (Burgess et al., 2018), FactorVAE (Kim and Mnih, 2018) and β TC-VAE (Chen et al., 2018). 2) Discrete VAEs, such as JointVAE (Dupont, 2018) and VQVAE (van den Oord et al., 2017). The JointVAE is a disentanglement method that jointly factorizes discrete and continuous variables, while VQVAE encodes discrete variables but does not encourage disentanglement. 3) Disentanglement models for text, such as Controlled Generation of Text (CGT) (Hu et al.,

	Continuous						Discrete		
	1	2	3	4	5	6	7	8	9
Verb tense	x	✓	x	✓	x	x	x	x	✓
Subj-num	x	✓	✓	✓	✓	✓	✓	✓	✓
Obj-num	x	✓	✓	✓	✓	✓	✓	✓	✓
Gender	✓	✓	✓	✓	x	x	x	✓	✓
Sent-type	✓	x	✓	✓	x	x	x	x	✓
Person-num	x	x	✓	x	x	x	x	x	✓
Verb-style	x	x	x	x	x	x	x	x	x
Negation	✓	✓	x	✓	✓	✓	x	✓	✓
Verb/obj	✓	✓	✓	✓	x	x	x	x	✓

TABLE 2: Summary of traversal for all latents. Models are abbreviated (1= β VAE, 2=CCIVAE, 3= β TCVAE, 4=FactorVAE, 5=CGT, 6=STVAE, 7=VQVAE, 8=JointVAE, 9=DCTC).

2017), which disentangles the sentence style, and Style Transfer VAE (ST-VAE) (John et al., 2019), that disentangles style from content.

Traversal Analysis. We consider a traversal for a factor to be disentangled only if the decoded sentences assume different values for the traversed factor, keeping all other factors unchanged, as explained in the previous section. The results for all generative factors are reported in Table 2, where the checkmark and cross symbolise respectively disentangled and entangled factors.

We see that DCTC achieves the most consistent semantics, by correctly disentangling 8 out of 9 generative factors. β TC-VAE and FactorVAE are able to disentangle respectively 6 and 7 out of 9 factors, while the NLP model (i.e. CGT and ST-VAE) are outperformed and only disentangle 3 out of 9 factors. We attribute the success of DCTC to the fact that it is modeling factors knowing their dimensions, and encourages disentanglement on the latents, after having encoded them as discrete variables. On the other hand, encoding discrete variables does not guarantee disentanglement by itself, as shown by the poor performance of VQ-VAE, while only accounting for style and content embeddings also leads to entangled representations, as shown by CGT and ST-VAE.

Samples of decoded sentences from the traversals for tense and subject-number are displayed in Table 3. DCTC correctly disentangles the tense into present, future and past, while the other factors are fixed, and similarly it disentangles the subject-number factor. On the other hand, β -VAE and JointVAE show an entangled representation for tense and subject-number.

	Tense	Subject-number
input	you will not attend the party	we will not attend the party
β VAE	you will not attend the party you will not sign the paper you will not attend the party	we will not attend the party he will not attend the party
JointVAE	you will not attend the party you did not join the wedding you do not attend the party	we will not attend the party you will not attend the party
DCTC	you will not attend the party you did not attend the party you do not attend the party	we will not attend the party i will not attend the party

TABLE 3: Traversal examples for tense and subject-number. Disentangled factors in blue, entangled in red.

5.2 Quantitative Evaluation

Measuring Text Disentanglement. Quantifying disentanglement is a necessary step in our evaluation, in order to make the qualitative assessment more granular. Current disentanglement metrics in Representation Learning (Higgins et al., 2016; Kim and Mnih, 2018), rely on image-based datasets that provide the true generative factors, however, most datasets do not present such annotations. Fortunately, text data has the advantage, over images, of being discrete and regular by nature, and thus, generative factors can be defined at a sentence-level, by leveraging linguistically robust syntactic and semantic categories. Following this intuition, a simple solution for being able to measure disentanglement in a text representation, is to have a pre-processing step, where generative factors are extracted, before utilizing this information to compute the quantitative metrics.

Data Preparation. In our experiment, we focus on the Yelp reviews dataset (Shen et al., 2017), which is composed by 600,000 review sentences, and we define and extract 5 generative factors, namely: gender, tense, negation, subject number, and object number. Using the part-of-speech (POS) engine provided by the Stanza python package (Qi et al., 2020), we extract: 1) the gender factor from the pronouns and 2) the number factor from subjects and objects. Similarly, the tense is obtained from the verb using Stanza’s lemmatizer, while negation is determined from the presence of negation attributes in the parsed metadata.

Experimental Setup. We follow the setup of the previous qualitative experiment in terms of baselines and parameters. We investigate the models on two datasets, namely, Yelp, where 5 generative factors are extracted with the described data prepa-

ration process, and dSentences, where 9 factors are provided. As a result, we set the DCTC model with respectively 5 and 9 latent variables, and in both cases we use the dimensions for linguistic features defined in Table 1.

Evaluation Metrics. We note that current disentanglement approaches in NLP (John et al., 2019) measure factors such as style transfer strength, content preservation, and quality of generation. In contrast, we are interested in evaluating the representations by computing the amount of disentanglement in the latents. To this end, we leverage the generative factors extracted from raw text, and compute three disentanglement metrics from the Representation Learning literature, namely *Z-diff* (Higgins et al., 2016), *Z-min-var* (Kim and Mnih, 2018), and Mutual Information Gap (MIG) (Chen et al., 2018). The main difference between *Z-diff*, *Z-min-var* and MIG is that the first two are reporting the accuracy of a classifier trained to recognize disentangled factors, while the last one is centered on measuring mutual information between latents and observed variables. More details about the metrics are reported in the appendix.

Metrics Analysis. The quantitative results in Table 4 show that DCTC outperforms the other baselines in 2 out of 3 metrics, for both datasets. Specifically, DCTC achieves the best performance in terms of MIG and *Z-min-var*, while obtaining the second best scores for the *Z-diff* metric after JointVAE and FactorVAE, respectively on the Yelp and dSentences datasets. These measurements are overall confirming the hypothesis of the qualitative experiment, namely, that the proposed DCTC model is able to achieve a disentangled representation of language, by jointly optimizing independence of latents and accounting for the discrete nature of the data.

Another similarity with the qualitative experiment is that both VQVAE and the continuous NLP baselines (CGT and ST-VAE) are on average outperformed by other models. The outcome of VQVAE indicates that despite the fact that discrete variables can encode the dimensions of linguistic features, a model should also encourage independence of the variables, in order to achieve competitive scores in disentanglement metrics. Similarly, we hypothesize that both CGT and ST-VAE do not necessarily consider all the generative factors that the experiment is evaluating, in fact, CGT only aims

		dSentences			Yelp		
		Z-min	Z-diff	MIG	Z-min	Z-diff	MIG
Cont.	CCI-VAE	0.79	0.71	0.23	0.83	0.84	0.25
	β VAE	0.88	0.87	0.32	0.91	0.75	0.30
	β TC-VAE	0.92	0.90	0.27	0.91	0.92	0.29
	FactorVAE	0.91	0.92	0.18	0.92	0.92	0.27
	CGT	0.78	0.63	0.13	0.77	0.66	0.18
	ST-VAE	0.82	0.67	0.24	0.84	0.72	0.26
Disc.	VQVAE	0.77	0.74	0.27	0.75	0.76	0.17
	JointVAE	0.89	0.81	0.35	0.90	0.95	0.33
	DCTC	0.94	0.91	0.43	0.94	0.92	0.49

TABLE 4: Disentanglement metrics results.

to disentangle the style embedding of sentences, while ST-VAE focuses on disentangling style from content.

5.3 Text Style Transfer

Arithmetics for Latent Factors. In this experiment, we take inspiration from previous work from Mikolov et al. (2013), which showed that word embeddings can capture semantic relations via vector arithmetics, (for example, king - man + woman = queen). More specifically, we consider text generative factors (e.g. negation), and investigate sentence-level embeddings arithmetic in the task of text style transfer (extrinsic evaluation of the models).

The style transfer protocol of our experiment is performed as follows. We first select a factor (e.g. negation) and extract two lists of sentences containing two specific values (e.g. negative, and affirmative), that we name respectively s_n and s_a . The extraction is performed based on the generative factors that we obtained with the pre-processing from our previous experiment. We then compute the vector of the arithmetic difference between the latents of the two vectors, namely $v = s_a - s_n$. Intuitively, this operation removes the negative components from the positive vector. Finally, we consider a third vector of negative sentences, encode them to obtain the embedding, and sum the previously computed vector v . After decoding we expect the sentences to be characterised by an affirmative style.

Experimental Setup. We follow our quantitative experiment for the model setup, and datasets. We compare our model with 3 state-of-the-art style transfer models, including: iVAE (Fang et al., 2019), DAAE (Shen et al., 2020), ST-VAE, along with top performing models from the previous experiments, namely FactorVAE, β TC-VAE, JointVAE. In terms of evaluation metrics, we report the style transfer accuracy for each factor. The accu-

		Gender	Negation	Tense	Subj	Obj
dSentences						
Cont.	FactorVAE	0.85	0.78	0.71	0.81	0.76
	β TC-VAE	0.70	0.76	0.45	0.75	0.81
	ST-VAE	0.79	0.81	0.56	0.79	0.92
	iVAE	0.82	0.85	0.56	0.90	0.88
	DAAE	0.89	0.93	0.56	0.89	0.85
	Disc.	JointVAE	0.72	0.83	0.67	0.87
	DCTC	0.90	0.94	0.73	0.86	0.95
Yelp						
Cont.	FactorVAE	0.83	0.89	0.23	0.72	0.80
	β TC-VAE	0.67	0.72	0.47	0.81	0.78
	ST-VAE	0.71	0.93	0.43	0.83	0.89
	iVAE	0.85	0.82	0.61	0.79	0.90
	DAAE	0.83	0.92	0.52	0.89	0.92
	Disc.	JointVAE	0.81	0.80	0.59	0.80
	DCTC	0.89	0.96	0.65	0.87	0.87

TABLE 5: Style Transfer Accuracy.

racy is computed by extracting the factors using the same procedure used for their selection.

Style Transfer Analysis. The style transfer accuracy results reported in Table 5 shows that our DCTC model outperforms other baselines for the majority of the factors, for both Yelp and dSentences. DCTC achieves achieves the second best result for subject-number after iVAE and DAAE, respectively on dSentences and Yelp.

We hypothesize that DCTC achieves the strongest performance due to the fact that it is explicitly set to learn variables with a known dimension, which can not be achieved with continuous models. Furthermore, by disentangling the latent variables, DCTC is able to provide a representation that results suitable for the task of flipping single factors. Finally, the strength of disentanglement is also highlighted by the fact that β TC-VAE and FactorVAE are performing comparably with the style transfer models, even if they are not created originally for this task.

Some examples for style transferred sentences from Yelp are reported in Table 6. We can observe that some baselines are able to invert the considered factors (tense and negation), however, DCTC is the only one that correctly inverts both factors without the need of changing other words from the input sentence. This can be justified by the ability of disentangled representation to encode invariance of certain factors. This extrinsic evaluation confirms the hypothesis that the joining disentangled representation with discrete encoding can positively impact the downstream task and represents a fundamental tool to design more expressive language encodings.

input	the pizza served was missing a portion
DAAE	the pizza is served a little bland
iVAE	the pizza served is for a ridiculous price
β VAE	the pizza was mistaken
β TC-VAE	the pizza is completely wrong!
ST-VAE	the pizza served is pretty decent
JointVAE	our pizza was great
DCTC	the pizza served is missing a portion

input	he told me he could not exchange them
DAAE	he told me he could exchange them
iVAE	he told me he could be more attentive
β VAE	he told me that he does not exchange them
β TC-VAE	he told me he wanted a second opinion
ST-VAE	he told me to try the plantains
JointVAE	he told us he could not be happier
DCTC	he told me he could exchange them

TABLE 6: Style transfer on Yelp. Top: tense, past to present. Bottom: negation, negative to affirmative. Correct changes in **blue**, wrong ones in **red**.

6 Conclusion

In this work, we propose the first approach where a discrete encoding of the linguistic features in sentences is integrated with an objective function that encourages disentanglement. We provide a VAE-based architecture where latent variables are back-propagated with a discrete reparameterization mechanism. We then design a decomposition of the ELBO, where 1) the independence between latent variables is encouraged, to aid disentanglement, and 2) the amount of encoded information is controlled, to avoid the posterior collapse. We provide a novel evaluation procedure where representations learned from text data are probed in terms of their disentanglement, using metrics from Representation Learning. With this evaluation tool, we demonstrate that the presented model consistently outperforms continuous and discrete baselines for disentanglement, on qualitative evaluation, quantitative metrics, and text style transfer.

We conclude that the modeling of discrete variables, which is currently under-explored in disentanglement research, may represent a fundamental encoding tool for enhancing interpretability and control in NLP models.

Acknowledgements

The research was partially funded by EPSRC and the BBC under iCASE and EPSRC Enncore (EP/T026995/1). The authors would like to thank Marco Valentino and Mokanarangan Thayaparan for the helpful discussions and Chris Newell for his project support.

References

- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. Plato: Pre-trained dialogue generation model with discrete latent variable. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 85–96.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- Samuel Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21.
- Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. 2018. Understanding disentangling in β -vae. *arXiv preprint arXiv:1804.03599*.
- Ricky TQ Chen, Xuechen Li, Roger Grosse, and David Duvenaud. 2018. Isolating sources of disentanglement in vaes. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 2615–2625.
- Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in neural information processing systems*, 29:2172–2180.
- Pengyu Cheng, Martin Renqiang Min, Dinghan Shen, Christopher Malon, Yizhe Zhang, Yitong Li, and Lawrence Carin. 2020. Improving disentangled text representation learning with information-theoretic guidance. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7530–7541.
- Pierre Colombo, Pablo Piantanida, and Chloé Clavel. 2021. A novel estimator of mutual information for learning to disentangle textual representations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, pages 6539–6550.
- Emilien Dupont. 2018. Learning disentangled joint continuous and discrete representations. In *Advances in Neural Information Processing Systems*, pages 710–720.
- Le Fang, Chunyuan Li, Jianfeng Gao, Wen Dong, and Changyou Chen. 2019. Implicit deep latent variable models for text generation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3937–3947.
- Daya Guo, Duyu Tang, Nan Duan, Jian Yin, Daxin Jiang, and Ming Zhou. 2020. Evidence-aware inferential text generation with vector quantised variational autoencoder. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6118–6129.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2016. beta-vae: Learning basic visual concepts with a constrained variational framework.
- Matthew D Hoffman and Matthew J Johnson. 2016. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Advances in Approximate Bayesian Inference, NIPS*, volume 1, page 2.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. 2017. Toward controlled generation of text. In *International Conference on Machine Learning*, pages 1587–1596. PMLR.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. Disentangled representation learning for non-parallel text style transfer. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 424–434.
- Hyunjik Kim and Andriy Mnih. 2018. Disentangling by factorising. In *International Conference on Machine Learning*, pages 2649–2658.
- Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR*.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. 2019. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pages 4114–4124. PMLR.
- Amine M’Charrak. 2018. Deep learning for natural language processing (nlp) using variational autoencoders (vae).
- Tomáš Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pages 746–751.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108.

- Lei Sha and Thomas Lukasiewicz. 2021. Multi-type disentanglement without adversarial training. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6833–6844.
- Tianxiao Shen, Jonas Mueller, Regina Barzilay, and Tommi Jaakkola. 2020. Educating text autoencoders: Latent representation guidance via denoising. In *International Conference on Machine Learning*, pages 8719–8729. PMLR.
- Lei Shu, Alexandros Papangelis, Yi-Chia Wang, Gokhan Tur, Hu Xu, Zhaleh Feizollahi, Bing Liu, and Piero Molino. 2020. Controllable text generation with focused variation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 3805–3817.
- Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. 2017. Neural discrete representation learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6309–6318.
- Satosi Watanabe. 1960. Information theoretical analysis of multivariate correlation. *IBM Journal of research and development*, 4(1):66–82.
- Shengjia Zhao, Jiaming Song, and Stefano Ermon. 2019. Infovae: Balancing learning and inference in variational autoencoders. In *Proceedings of the aai conference on artificial intelligence*, volume 33, pages 5885–5892.

A Disentanglement Metrics

Z-diff Z-diff (Higgins et al., 2016) considers pairs of instances to create batches where a generative factor k is chosen randomly. We then consider the pairs $v_1 v_2$ that have the same value for the factor k . The absolute difference of the encoding of the pair is then determined, namely $|z_1 - z_2|$. The intuition of the metric is that a smaller difference for a fixed factor entails more similar samples. A reference dataset is constructed with this procedure, then a linear classifier is trained to predict which factor is fixed, and the accuracy is considered as the disentanglement metric.

Z-min-var The Z-min-var (Kim and Mnih, 2018) is similar to the Z-diff metric, as it creates a reference dataset and train a classifier to find the fixed factor. Specifically, Z-min-var builds the dataset using the *argmin* of the variance vector of all encodings, which have been normalized by the standard deviation.

Z-diff and Z-min-var both rely on the intuition that a smaller difference for a fixed factor entails more similar samples. They create a dataset for a classifier to predict the fixed factor, the accuracy of which is a measure for disentanglement. Z-diff creates the dataset with the absolute differences of encoding pairs where a factor is fixed, while Z-min-var builds a similar dataset, but using the *argmin* of the variance vector of all encodings, which have been normalized by the standard deviation.

Mutual Information Gap MIG (Chen et al., 2018) does not rely on a classifier, and thus it provides more robustness against hyperparameter biases. MIG first computes the mutual information between each latent and the true factor, and then it identifies and subtracts the two values for latents with maximum mutual information. The obtained quantity is considered the amount of disentanglement.