

MITIGATING POSTERIOR COLLAPSE IN STRONGLY CONDITIONED VARIATIONAL AUTOENCODERS

Anonymous authors

Paper under double-blind review

ABSTRACT

Training conditional generative latent-variable models is challenging in scenarios where the conditioning signal is very strong and the decoder is expressive enough to generate a plausible output given only the condition; the generative model tends to ignore the latent variable, suffering from posterior collapse. We find, and empirically show, that one of the major reasons behind posterior collapse is rooted in the way that generative models are conditioned, i.e., through concatenation of the latent variable and the condition. To mitigate this problem, we propose to explicitly make the latent variables depend on the condition by unifying the conditioning and latent variable sampling, thus coupling them so as to prevent the model from discarding the root of variations. To achieve this, we develop a conditional Variational Autoencoder architecture that learns a distribution not only of the latent variables, but also of the condition, the latter acting as prior on the former. Our experiments on the challenging tasks of conditional human motion prediction and image captioning demonstrate the effectiveness of our approach at avoiding posterior collapse. Video results of our approach are anonymously provided¹ in <http://bit.ly/iclr2020>.

1 INTRODUCTION

Deep generative models offer promising results in generating diverse, realistic samples, such as images, text, motion, and sound, from purely unlabeled data. One example of such successful generative models are variational autoencoders (Kingma & Welling, 2013) (VAEs), the stochastic variant of autoencoders, which, thanks to strong and expressive decoders, can generate high quality samples. Training such models, however, may often result in posterior collapse: the posterior distribution $q(z|x)$ of the latent variable z given the input x becomes equal to the prior distribution, resulting in a latent variable carrying no information about the input. In other words, the model learns to ignore the latent variable.

The most common approaches to tackling posterior collapse consist of weighing the KL divergence between the posterior and prior during training by an annealing function (Bowman et al., 2015; Yang et al., 2017; Kim et al., 2018; Gulrajani et al., 2016; Liu et al., 2019), weakening the decoder (Semeniuta et al., 2017; Zhao et al., 2017), or changing the training objective (Zhao et al., 2017; Tolstikhin et al., 2017). All of them are based on the perspective that the solution to posterior collapse can be found in a good local optimum in terms of evidence lower bound (Chen et al., 2016; Alemi et al., 2017). However, they each suffer from drawbacks: Any annealing weight that does not become, and remain equal to one at some point during training yields an improper statistical model; weakening the decoder tends to degrade the quality of the generated samples; changing the objective does not optimize the true variational lower bound. As an alternative, some methods modify the training strategy to more strongly encourage the inference network to approximate the model’s true posterior (He et al., 2019; Li et al., 2019). Other methods add auxiliary tasks either with non-autoregressive models (Lucas & Verbeek, 2018) or that exploit the latent variable (Goyal et al., 2017; Lucas & Verbeek, 2018; Dieng et al., 2018). While this encourages the latent variable to carry some information, it may not be directly useful for the main task. Alternatively, several techniques incorporate constraints in VAEs. In this context, VQ-VAE (van den Oord et al., 2017) introduces a discrete latent variable obtained by vector quantization of the latent one that, given a uniform prior

¹Video results are anonymously shared with ICLR 2020 Program Chair’s approval.

over the outcome, yields a fixed KL divergence equal to $\log K$, with K the size of the codebook; several recent works use the von Mises-Fisher distribution to obtain a fixed KL divergence, thus mitigating the posterior collapse problem (Guu et al., 2018; Xu & Durrett, 2018; Davidson et al., 2018); more recently, delta-VAE (Razavi et al., 2019) modifies the posterior such that it maintains a minimum distance between the prior and the posterior.

Although successful at handling the posterior collapse in presence of expressive decoders, e.g., LSTM, GRUs, PixelCNN, all of these approaches were designed for standard VAEs, not *conditional* VAEs (CVAEs). As such, they do not address the problem of mitigating the influence of a strong conditioning signal in ignoring the latent variable, which is our focus here. In this context, we observed that a strong condition provides enough information for an expressive decoder to reconstruct the data, thus allowing the decoder to ignore the latent variable at no loss in reconstruction quality.

To address this problem, in this paper, we exploit the following observation: In standard CVAEs, the random variable is sampled from the prior/posterior completely independently of the conditioning one. This decouples the two sources of information, i.e., the condition and the random latent variable, and opens the door for the model to discard the latter. Therefore, here, we propose to explicitly make the random variable depend on the conditioning one. To this end, we develop a CVAE architecture that learns a distribution not only of the latent variable but also of the conditioning one. We then use this distribution as a prior over the latent variable, making its sampling explicitly dependent on the condition. As such, we name our method **CPP-VAE**, for **Condition Posterior as Prior**.

We empirically show the effectiveness of our approach for problems that are stochastic in nature. In particular, we focus on scenarios where the training dataset is deterministic, i.e., one condition per data sample, and the conditioning signal strong enough for an expressive decoder to generate a plausible output from it. This, for example, occurs in human motion prediction, that is, forecasting future 3D poses given a sequence of observed ones. In this context, existing methods typically fail to model the stochastic nature of human motion, either because they learn a deterministic mapping from the observations to the output, or because the stochastic latent variables they combine with the observations can be ignored by the model. As an alternative application, we also evaluate our approach on image captioning, i.e., generating diverse and plausible captions describing an image. Our empirical results show that not only does our approach yield a much wider variety of plausible samples than concatenation-based stochastic methods, but it also preserves the semantic information of the condition, such as the type of action performed by the person in motion prediction or visual image elements in captioning, without explicitly exploiting this information. We will make the source code for our model and for our evaluation metrics publicly available upon acceptance.

2 BACKGROUND

In this section, we briefly review the basics of VAEs and conditional VAEs that will be necessary to introduce our method in the next section. We then briefly introduce the posterior collapse phenomenon in the context of conditional VAEs.

Variational Autoencoders. VAEs are a family of generative models that utilize neural networks to learn the distribution of the data. To this end, VAEs first learn to generate a latent variable z given the data x , i.e., approximate the posterior distribution $q_\phi(z|x)$, where ϕ are the parameters of a neural network, the encoder, whose goal is to model the variation of the data. From this latent random variable z , VAEs then generate a sample x by learning $p_\theta(x|z)$, where θ denotes the parameters of another neural network, the decoder, whose goal is to maximize the log likelihood of the data.

These two networks, i.e., the encoder ($q_\phi(z|x)$) and the decoder ($p_\theta(x|z)$), are trained jointly, using a prior over the latent variable. This prior is usually the standard Normal distribution, $\mathcal{N}(0, I)$. Note that VAEs use a variational approximation of the posterior, i.e., $q_\phi(z|x)$, rather than the true posterior. This enables the model to maximize the variational lower bound of the log likelihood with respect to the parameters ϕ and θ , given by

$$\log p_\theta(x) \geq \mathbb{E}_{q_\phi(z|x)} \left[\log p_\theta(x|z) \right] - KL \left(q_\phi(z|x) || p(z) \right), \quad (1)$$

where the second term encodes the KL divergence between the posterior and the prior distributions. In practice, the posterior distribution is approximated by a Gaussian $\mathcal{N}(\mu, \text{diag}(\sigma^2))$, whose parameters are output by the encoder. Note, σ is a vector and we define σ^2 as a vector whose elements

are the squared elements of σ . To facilitate optimization, the reparameterization trick (Kingma & Welling, 2013) is used. That is, the latent variable is computed as $z = \mu + \sigma \odot \epsilon$, where ϵ is a vector sampled from the standard Normal distribution.

Conditional Variational Autoencoders. As an extension to VAEs, CVAEs use auxiliary information, i.e., the conditioning variable or observation, to generate the data x . In the standard setting, both the encoder and the decoder are conditioned on the conditioning variable c . That is, the encoder is denoted as $q_\phi(z|x, c)$ and the decoder as $p_\theta(x|z, c)$. Then, the objective of the model becomes

$$\log p_\theta(x|c) \geq \mathbb{E}_{q_\phi(z|x, c)} \left[\log p_\theta(x|z, c) \right] - KL\left(q_\phi(z|x, c) || p(z|c)\right). \quad (2)$$

In practice, conditioning is typically done by concatenation; the input of the encoder is the concatenation of the data x and the condition c , i.e., $q_\phi(z|[x, c])$, and that of the decoder the concatenation of the latent variable z and the condition c , i.e., $p_\theta(x|[z, c])$. Thus, the prior distribution is still $p(z)$, and the latent variable is sampled independently of the conditioning one. It is then left to the decoder to combine the information from the latent and conditioning variables to generate a data sample. The detailed computations of the VAE/CVAE lower bounds are provided in Appendix A.

Posterior Collapse. As mentioned before, training conditional generative latent-variable models is challenging due to posterior collapse. This can be observed when the KL divergence term becomes zero, which means that, regardless of the input, the approximate posterior distribution is equal to the prior distribution. In other words, there is no semantic connection between the encoder and the decoder, and thus the latent variable drawn from the approximate posterior does not convey any useful information to obtain an input-dependent reconstruction. In this case, the decoder generates samples that approximate the mean of the whole training set, minimizing the reconstruction loss. We found that one of the major reasons behind posterior collapse in the case of conditional VAEs with strong conditioning signals and expressive decoders is rooted in the conventional way of conditioning, i.e., through concatenation of the latent variable and the condition. Concatenation allows the decoder to decouple the latent variable from the deterministic condition, thus allowing the decoder to optimize its reconstruction loss given only the condition. In the following, we introduce our approach for unifying latent variable sampling and conditioning, so that the decoder cannot decouple the two sources of information.

3 UNIFYING SAMPLING LATENT VARIABLE AND CONDITIONING

In this section, we introduce our approach as a general framework to mitigate posterior collapse and thus generating diverse and plausible samples using a CVAE. In essence, our framework consists of two autoencoders, one acting on the conditioning signal and the other on the samples we wish to learn the distribution of. The latent representation of the condition then serves as conditioning variable to generate the desired samples.

As discussed above, we are interested in problems that are stochastic in nature; given one condition, multiple plausible and natural samples are likely. However our training data is insufficiently sampled, in that for any given condition, the dataset contains only a single observed sample, in effect making the data appear deterministic. Moreover, in these cases, the condition provides the core signal to generate a good sample, even in a deterministic model. Therefore, it is highly likely that a CVAE trained for this task learns to ignore the latent variable and rely only on the condition to produce its output. Below, we address this by forcing the sampling of the random latent variable to depend on the conditioning one. By making this dependency explicit, we prevent the network from ignoring the latent variable in the presence of a strong condition, thus enabling it to generate diverse outputs.

Note that conditioning the VAE *encoder* via standard strategies, e.g., concatenation, is perfectly fine, since the two inputs to the encoder are deterministic and useful to compress the sample into the latent space. However, conditioning the VAE *decoder* requires special care to avoid posterior collapse and make sure that the latent space carries some information. This is what we focus on below.

Stochastically Conditioning the Decoder. To avoid posterior collapse, we propose to make the sampling of the latent variable from the prior/posterior distribution explicitly depend on the condi-

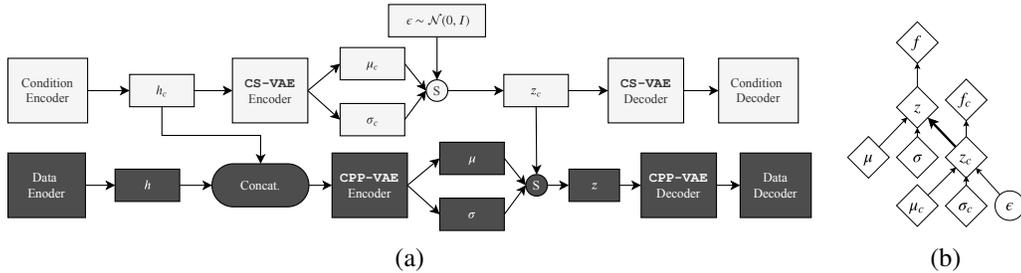


Figure 1: (a) Overview of our CPP-VAE framework. (b) Our extension to the reparameterization trick for sampling a latent variable given the condition. A circle denotes a random node, a diamond denotes a deterministic node, and f is the objective. In our design, $\epsilon \sim \mathcal{N}(0, I)$.

tion instead of treating these two variables as independent. To this end, we first learn the distribution of the condition via a simple VAE, which we refer to as CS-VAE because this VAE acts on the conditioning signal. The goal of CS-VAE is to reconstruct the condition given its latent representation. We take the prior of CS-VAE as a standard Normal distribution $\mathcal{N}(0, I)$. Following (Kingma & Welling, 2013), this allows us to approximate the CS-VAE posterior with another sample from a Normal distribution $\epsilon \sim \mathcal{N}(0, I)$. That is, we write

$$z_c = \mu_c + \sigma_c \odot \epsilon, \quad (3)$$

where μ_c and σ_c are the parameters of the posterior distribution generated by the VAE encoder.

Following the same strategy for the VAE on the data, called CPP-VAE, translates to treating the conditioning and the data latent variables independently, which we want to avoid. Therefore, as illustrated in Fig. 1 (Right), we instead define the CPP-VAE posterior as not directly normally distributed, but conditioned on the posterior of CS-VAE. To this end, we extend the standard reparameterization as

$$z = \mu + \sigma \odot z_c = \mu + \sigma \odot (\mu_c + \sigma_c \odot \epsilon) = \underbrace{(\mu + \sigma \odot \mu_c)}_{\text{CPP-VAE's mean}} + \underbrace{(\sigma \odot \sigma_c)}_{\text{CPP-VAE's std.}} \odot \epsilon, \quad (4)$$

where z_c comes from Eq. 3. In fact, z_c in Eq. 3 is a sample from the scaled and translated version of $\mathcal{N}(0, I)$ given μ_c and σ_c , and z in Eq. 4 is a sample from the scaled and translated version of $\mathcal{N}(\mu_c, \text{diag}(\sigma_c^2))$ given μ and σ . Since we have access to the observations during both training and testing, we always sample z_c from the condition posterior. As z is sampled given z_c , one expects the latent variable z to carry information about the strong condition, and thus a sample generated from z to correspond to a plausible sample given the condition. This extended reparameterization trick allows us to avoid conditioning the CPP-VAE decoder by concatenating the latent variable with a deterministic representation of the condition, thus mitigating posterior collapse. However, it changes the variational family of the CPP-VAE posterior. In fact, the posterior is no longer $\mathcal{N}(\mu, \text{diag}(\sigma^2))$, but a Gaussian distribution with mean $\mu + \sigma \odot \mu_c$ and covariance matrix $\text{diag}((\sigma \odot \sigma_c)^2)$. This will be accounted for when designing the KL divergence loss discussed below.

Learning. To learn the parameters of our model, we rely on the availability of a dataset $D = \{X_1, X_2, \dots, X_N\}$ containing N training samples X_i . Each training sample is a pair of condition and desired sample. For CS-VAE, that learns the distribution of the condition, we define the loss as the KL divergence between its posterior and the standard Gaussian prior, that is,

$$\mathcal{L}_{\text{prior}}^{\text{CS-VAE}} = KL\left(\mathcal{N}(\mu_c, \text{diag}(\sigma_c^2)) \parallel \mathcal{N}(0, I)\right) = -\frac{1}{2} \sum_j \left(1 + \log(\sigma_{c_j}^2) - \mu_{c_j}^2 - \sigma_{c_j}^2\right). \quad (5)$$

By contrast, for CPP-VAE, we define the loss as the KL divergence between the posterior of CPP-VAE and the posterior of the CS-VAE, i.e., of the condition. To this end, we freeze the weights of CS-VAE before computing the KL divergence, since we do not want to move the posterior of the condition but that of the data. The KL divergence is then computed as the divergence between two multivariate Normal distributions, encoded by their mean vectors and covariance matrices, as

$$\mathcal{L}_{\text{prior}}^{\text{CPP-VAE}} = KL\left(\mathcal{N}(\mu + \sigma \odot \mu_c, \text{diag}((\sigma \odot \sigma_c)^2)) \parallel \mathcal{N}(\mu_c, \text{diag}(\sigma_c^2))\right). \quad (6)$$

Let $\Sigma = \text{diag}(\sigma^2)$, $\Sigma_c = \text{diag}(\sigma_c^2)$, d be the dimensionality of the latent space and $\text{tr}\{\cdot\}$ the trace of a square matrix, the loss in Eq. (6) can be written as²

$$\mathcal{L}_{prior}^{CPP-VAE} = -\frac{1}{2} \left[\log \frac{1}{|\Sigma|} - d + \text{tr}\{\Sigma\} + (\mu_c - (\mu + \Sigma\mu_c))^T \Sigma_c^{-1} (\mu_c - (\mu + \Sigma\mu_c)) \right]. \quad (7)$$

After computing the loss in Eq. 7, we unfreeze CS-VAE and update it with its previous gradient. Trying to match the posterior of CPP-VAE to that of CS-VAE allows us to effectively use our extended reparameterization trick in Eq. 4. Furthermore, we use the standard reconstruction loss for both CS-VAE and CPP-VAE, minimizing the negative log-likelihood (NLL) or the mean squared error (MSE) of the condition and the corresponding data, given the task. We refer to the reconstruction losses as $\mathcal{L}_{rec}^{CS-VAE}$ and $\mathcal{L}_{rec}^{CPP-VAE}$ for CS-VAE and CPP-VAE. Thus, our complete loss is

$$\mathcal{L} = \lambda(\mathcal{L}_{prior}^{CS-VAE} + \mathcal{L}_{prior}^{CPP-VAE}) + \mathcal{L}_{rec}^{CS-VAE} + \mathcal{L}_{rec}^{CPP-VAE}. \quad (8)$$

In practice, since our VAE appears within a recurrent model, we weigh the KL divergence terms by a function λ corresponding to the KL annealing weight of (Bowman et al., 2015). We start from $\lambda = 0$, forcing the model to encode as much information in z as possible, and gradually increase it to $\lambda = 1$ during training, following a logistic curve. We then continue training with $\lambda = 1$.

In short, our method can be interpreted as a simple yet effective framework (designed for CVAEs) for altering the variational family of the posterior so as to prevent posterior collapse by making sure that there is a positive mismatch between the two distributions in the KL loss of Eq. 7.

4 EXPERIMENTS

In this paper, we mainly focus on stochastic human motion prediction, where given partial observation, the task is to generate diverse and plausible continuations. Additionally, to show that our CPP-VAE generalizes to other domains, we tackle the problem of stochastic image captioning, where given an image representation, the task is to generate diverse yet related captions.

4.1 DIVERSE HUMAN MOTION PREDICTION

Dataset. To evaluate the effectiveness of our approach on the task of stochastic human motion prediction, we use the Human3.6M dataset (Ionescu et al., 2014), the largest publicly-available motion capture (mocap) dataset. Human3.6M comprises more than 800 long indoor motion sequences performed by 11 subjects, leading to 3.6M frames. Each frame contains a person annotated with 3D joint positions and rotation matrices for all 32 joints. In our experiments, for our approach and the replicated VAE-based baselines, we represent each joint in 4D quaternion space. We follow the standard preprocessing and evaluation settings used in (Martinez et al., 2017; Gui et al., 2018a; Pavlo et al., 2018; Jain et al., 2016). We also evaluate our approach on a real-world dataset, Penn Action (Zhang et al., 2013), which contains 2326 sequences of 15 different actions, where for each person, 13 joints are annotated in 2D space. The results on Penn Action are provided in Appendix F.

Evaluation Metrics. To quantitatively evaluate our approach and other stochastic motion prediction baselines (Yan et al., 2018; Barsoum et al., 2018; Walker et al., 2017; Aliakbarian et al., 2019), we report the estimated upper bound on the reconstruction error as ELBO, along with the KL-divergence on the held-out test set. Additionally, we also use quality (Aliakbarian et al., 2019) and diversity (Yang et al., 2019; Aliakbarian et al., 2019; Yuan & Kitani, 2019) metrics (which should be considered together), a context metric, and the training KL at convergence. To measure the diversity of the motions generated by a stochastic model, we make use of the average distance between all pairs of the K motions generated from the same observation. To measure quality, we train a binary classifier to discriminate real (ground-truth) samples from fake (generated) ones. The accuracy of this classifier on the test set is inversely proportional to the quality of the generated motions. Context is measured by the performance of a good action classifier (Li et al., 2018) trained on ground-truth motions. The classifier is then tested on each of the K motions generated from each observation. For N observations and K continuations per observation, the accuracy is measured by computing

²See Appendix B for more details on the KL divergence between two multivariate Gaussians and the derivation of Eq. 7.

Table 1: Comparison of CPP-VAE with the stochastic motion prediction baselines on Human3.6M.

Method	ELBO (KL) (Reconstructed)	Diversity (Sampled)	Quality (Sampled)	Context (Sampled)	Training KL (Reconstructed)
MT-VAE (Yan et al., 2018)	0.51 (0.06)	0.26	0.45	0.42	0.08
Pose-Knows (Walker et al., 2017)	2.08 (N/A)	1.70	0.13	0.08	N/A
HP-GAN (Barsoum et al., 2018)	0.61 (N/A)	0.48	0.47	0.35	N/A
Mix-and-Match (Aliakbarian et al., 2019)	0.55 (2.03)	3.52	0.42	0.37	1.98
CPP-VAE	0.41 (8.07)	3.12	0.48	0.54	6.93

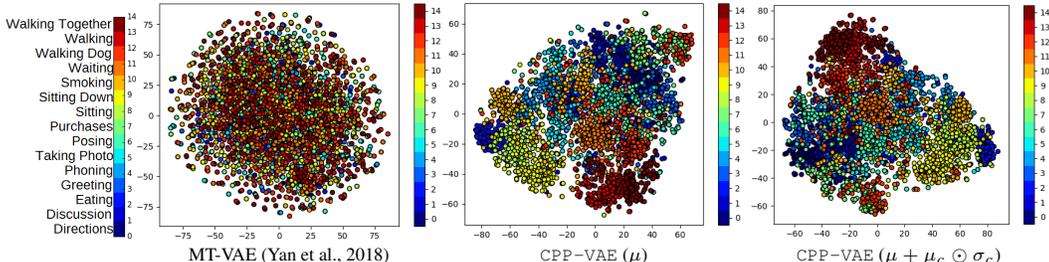


Figure 2: t-SNE plots of the posterior mean for 3750 test motions. With MT-VAE (Yan et al., 2018), all classes are mixed, suggesting that the latent variable carries little information about the input. By contrast, our condition-dependent sampling allows CPP-VAE to better preserve context. Note that some actions, such as discussion and directions, are very hard to identify and are thus spread over other actions. Others, such as walking, walking with dog, and walking together or sitting and sitting down overlap due to their similarity.

the argmax over each prediction’s probability vector, and we report context as the mean class accuracy on the $K \times N$ motions. For all metrics, we use $K = 50$ motions per test observation. We also provide qualitative results in Appendix L. For all experiments related to motion prediction, we use 16 frames (i.e., 640ms) as observation to generate the next 60 frames (i.e., 2.4sec).

Evaluating Stochasticity. In Table 1, we compare our approach (with the architecture described in Appendix H) with the state-of-the-art stochastic motion prediction models (Yan et al., 2018; Aliakbarian et al., 2019; Walker et al., 2017; Barsoum et al., 2018). Note that one should consider the reported metrics jointly to truly evaluate a stochastic model. For instance, while MT-VAE (Yan et al., 2018) and HP-GAN (Barsoum et al., 2018) generate high-quality motions, they are not diverse. Conversely, while Pose-Knows (Walker et al., 2017) generates diverse motions, they are of low quality. On the other hand, our approach generates both high quality and diverse motions. This is also the case of Mix-and-Match (Aliakbarian et al., 2019), which, however, preserves much less context. In fact, none of the baseline can effectively convey the context of the observation to the generated motions properly. As shown in Table 2, the upper bound for the context on Human3.6M is 0.60 (i.e., the classifier (Li et al., 2018) performance given the ground-truth motions). Our approach yields a context of 0.54 when given only about 20% of the data. Altogether, our approach yields diverse, high-quality and context-preserving predictions. This is further evidenced by the t-SNE (Maaten & Hinton, 2008) plots of Fig. 2, where different samples of various actions are better separated for our approach than for, e.g., MT-VAE (Yan et al., 2018). We refer the reader to the human motion prediction related work section in Appendix C for a brief overview of the baselines. We also encourage reading Appendix D for further discussion of the aforementioned baselines and a deeper insight of their behavior under different evaluation metrics.

Evaluating Sampling Quality. To further evaluate the sampling quality, we evaluate stochastic baselines using the standard mean angle error (MAE) metric in Euler space. To this end, we use the best of the $K = 50$ generated motions for each observation (aka S-MSE (Yan et al., 2018)). A model that generates more diverse motions has more chances of generating a motion close to the ground-truth one. As shown in Table 3, this is the case with our approach and Mix-and-Match (Aliakbarian

Table 2: Comparison of the generated motions with the ground-truth future motion in terms of context. The gap between the performance of the state-of-the-art pose-based action classifier (Li et al., 2018) with and without true future motions is 0.22. Using our predictions, this gap decreases to 0.06, showing that our predictions reflect the class label.

Note	Observation	Future Motion	Context
Lower bound (no additional info to the observation)	ground-truth	Zero velocity	0.38
Upper bound (ground-truth poses as future motion)	ground-truth	ground-truth	0.60
Ours (sampled motions as future motion)	ground-truth	Sampled from CPP-VAE	0.54

Table 3: Comparison with the state-of-the-art stochastic motion prediction models for 4 actions of Human3.6M (all methods use the best of $K = 50$ sampled motions).

Method	Walking					Eating					Smoking					Discussion								
	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
MT-VAE (Yan et al., 2018)	0.73	0.79	0.90	0.93	0.95	1.05	0.68	0.74	0.95	1.00	1.03	1.38	1.00	1.14	1.43	1.44	1.68	1.99	0.80	1.01	1.22	1.35	1.56	1.69
HP-GAN (Barsoum et al., 2018)	0.61	0.62	0.71	0.79	0.83	1.07	0.53	0.67	0.79	0.88	0.97	1.12	0.64	0.78	1.05	1.12	1.64	1.84	0.79	1.00	1.12	1.29	1.43	1.71
Pose-Knows (Walker et al., 2017)	0.56	0.66	0.98	1.05	1.28	1.60	0.44	0.60	0.71	0.84	1.05	1.54	0.59	0.83	1.25	1.36	1.67	2.03	0.73	1.10	1.33	1.34	1.45	1.85
Mix&Match (Aliakbarian et al., 2019)	0.33	0.48	0.56	0.58	0.64	0.68	0.23	0.34	0.41	0.50	0.61	0.91	0.23	0.42	0.79	0.77	0.82	1.25	0.25	0.60	0.83	0.89	1.12	1.30
CPP-VAE	0.22	0.36	0.47	0.52	0.58	0.69	0.19	0.28	0.40	0.51	0.58	0.90	0.23	0.43	0.77	0.75	0.78	1.23	0.21	0.52	0.81	0.84	1.04	1.28

Table 4: Comparison with the state-of-the-art deterministic models for 4 actions of Human3.6M.

Method	Walking					Eating					Smoking					Discussion								
	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000	80	160	320	400	560	1000
Zero Velocity	0.39	0.86	0.99	1.15	1.35	1.32	0.27	0.48	0.73	0.86	1.04	1.38	0.26	0.48	0.97	0.95	1.02	1.69	0.31	0.67	0.94	1.04	1.41	1.96
LSTM-3LR (Fragkiadaki et al., 2015)	1.18	1.50	1.67	1.76	1.81	2.20	1.36	1.79	2.29	2.42	2.49	2.82	2.05	2.34	3.10	3.18	3.24	3.42	2.25	2.33	2.45	2.46	2.48	2.93
SRNN (Jain et al., 2016)	1.08	1.34	1.60	1.80	1.90	2.13	1.35	1.71	2.12	2.21	2.28	2.58	1.90	2.30	2.90	3.10	3.21	3.23	1.67	2.03	2.20	2.31	2.39	2.43
DAE-LSTM (Ghosh et al., 2017)	1.00	1.11	1.39	1.48	1.55	1.39	1.31	1.49	1.86	1.89	1.76	2.01	0.92	1.03	1.15	1.25	1.38	1.77	1.11	1.20	1.38	1.42	1.53	1.73
GRU (Martinez et al., 2017)	0.28	0.49	0.72	0.81	0.93	1.03	0.23	0.39	0.62	0.76	0.95	1.08	0.33	0.61	1.05	1.15	1.25	1.50	0.31	0.68	1.01	1.09	1.43	1.69
AGED (Gui et al., 2018a)	0.22	0.36	0.55	0.67	0.78	0.91	0.17	0.28	0.51	0.64	0.86	0.93	0.27	0.43	0.82	0.84	1.06	1.21	0.27	0.56	0.76	0.83	1.25	1.30
DCT-GCN (Mao et al., 2019)	0.18	0.31	0.49	0.56	0.65	0.67	0.16	0.29	0.50	0.62	0.76	1.12	0.22	0.41	0.86	0.80	0.87	1.57	0.20	0.51	0.77	0.85	1.33	1.70
CPP-VAE ($z = \mu_c$)	0.20	0.34	0.48	0.53	0.57	0.71	0.20	0.26	0.44	0.52	0.61	0.92	0.21	0.43	0.79	0.79	0.77	1.15	0.22	0.55	0.79	0.81	1.05	1.28

et al., 2019), which both yield higher diversity. However, our approach performs better thanks to its context-preserving latent representation and its higher quality of the generated motions.

In Table 4, we compare our approach with the state-of-the-art deterministic motion prediction models (Martinez et al., 2017; Jain et al., 2016; Gui et al., 2018b; Fragkiadaki et al., 2015; Gui et al., 2018a) using the MAE metric in Euler space. To have a fair comparison, we generate one motion per observation by setting the latent variable to the distribution mode, i.e., $z = \mu_c$. This allows us to generate a plausible motion without having access to the ground-truth. To compare against the deterministic baselines, we follow the standard setting, and thus use 50 frames (i.e., 2sec) as observation to generate the next 25 frames (i.e., 1sec). Surprisingly, despite having a very simple motion decoder architecture (one-layer GRU network) with a very simple reconstruction loss function (MSE), this motion-from-mode strategy yields results that are competitive with those of the baselines that use sophisticated architectures and advanced loss functions. We argue that learning a good, context-preserving latent representation of human motion is the contributing factor to the success of our approach. This, however, could be used in conjunction with sophisticated motion decoders and reconstruction losses, which we leave for future research.

In Appendix E, we study alternative designs to condition the VAE encoder and decoder.

4.2 DIVERSE IMAGE CAPTIONING

For the task of conditional text generation, we focus on stochastic image captioning. To demonstrate the effectiveness of our approach, we report results on the MSCOCO (Lin et al., 2014) captioning task with the original train/test splits of 83K and 41K images, respectively. The MSCOCO dataset has five captions per image. However, we make it deterministic by removing four captions per image, yielding a Deterministic-MSCOCO captioning dataset. Note that the goal of this experiment is not to advance the state of the art in image captioning, but rather to explore the effectiveness of our approach on a different task, where we have strong conditioning signal and an expressive decoder in the presence of a deterministic dataset.

Table 5: Quantitative evaluation of stochastic image captioning on the MSCOCO Captioning dataset.

Model	ELBO (KL) (Reconstructed)	Perplexity (Reconstructed)	Quality (Sampled)	Diversity (Sampled)	Context (Sampled)	Training KL (Reconstructed)
Autoregressive	3.01 (N/A)	20.29	0.40	N/A	0.46	N/A
Conditional VAE	2.86 (0.00)	17.46	0.39	0.00	0.44	0.00
CPP-VAE	0.21 (3.28)	1.23	0.40	0.53	0.43	3.11

A brief review of the recent work on diverse text generation is given in Appendix J.

We compare CPP-VAE (with the architecture described in Appendix I) with a standard CVAE and with its autoregressive, non-variational counterpart³. For quantitative evaluation, we report the ELBO (the negative log-likelihood), along with the KL-divergence and the Perplexity of the reconstructed captions on the held-out test set. We also quantitatively measure the diversity, the quality, and the context of sampled captions. To measure the context, we rely on the BLEU1 score, making sure that the sampled captions represent elements that appear in the image. For CVAE and CPP-VAE, we compute the average BLEU1 score for $K = 50$ captions sampled per image and report the mean over the images. To measure the diversity, we measure the BLEU4 score between every pair of $K = 50$ sampled captions per image. The smaller the BLEU4 is, the more diverse the captions are. The diversity metric is then $1 - \text{BLEU4}$, i.e., the higher the better. To measure the quality, we use a metric similar to that in our human motion prediction experiments, obtained by training a binary classifier to discriminate real (ground-truth) captions from fake (generated) ones. The accuracy of this classifier on the test set is inversely proportional to the quality of the generated captions. We expect a good stochastic model to have high quality and high diversity at the same time, while capturing the context of the given image. We provide qualitative examples for all the methods in Appendix M. As shown in Table 5, a CVAE learns to ignore the latent variable as it can minimize the caption reconstruction loss given solely the image representation. By doing so, all the generated captions at test time are identical, despite sampling multiple latent variables. This can be further seen in the ELBO and Perplexity of the reconstructed captions. We expect a model that gets as input the captions and the image to have a much lower reconstruction loss compared to the autoregressive baseline (which gets only the image as input). However, this is not the case with CVAE, indicating that the connection between the encoder and the decoder, i.e., the latent variable, does not carry essential information about the input caption. However, the quality of the generated sample is reasonably good. This is also illustrated in the qualitative evaluations in Appendix M. CPP-VAE, on the other hand, is able to effectively handle this situation by unifying the sampling of the latent variable and the conditioning, leading to diverse but high quality captions, as reflected by the ELBO of our approach in Table 5 and the qualitative results in Appendix M. Additional quantitative evaluations and ablation studies for image captioning are provided in Appendix K.

5 CONCLUSION

In this paper, we have studied the problem of posterior collapse in conditional generative latent-variable models, with a focus on scenarios where the conditioning signal is strong enough such that an expressive decoder can generate plausible samples from it only. We have addressed this problem by forcing the sampling of the random latent variable to depend on the conditioning one. By making this dependency explicit, we have prevented the network from ignoring the latent variable in the presence of a strong condition, thus enabling it to generate diverse outputs. To demonstrate the effectiveness of our approach, we have investigated two application domains: Stochastic human motion prediction and diverse image captioning. In both cases, our CPP-VAE model was able to generate diverse and plausible samples, as well as to retain contextual information, leading to semantically-meaningful predictions. In the future, we will apply our approach to other problems that rely on strong conditions, such as image inpainting and super-resolution, for which only deterministic datasets are available.

³Note that CPP-VAE is agnostic to the choice of data encoder/decoder architecture. Thus, one could use more sophisticated architectures, which we leave for future research.

REFERENCES

- Alexander A Alemi, Ben Poole, Ian Fischer, Joshua V Dillon, Rif A Saurous, and Kevin Murphy. Fixing a broken elbow. *arXiv preprint arXiv:1711.00464*, 2017.
- Mohammad Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Basura Fernando, Lars Petersson, and Lars Andersson. Viena²: A driving anticipation dataset. In *Asian Conference on Computer Vision*, pp. 449–466. Springer, 2018.
- Mohammad Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, Stephen Gould, and Amirhossein Habibian. Learning variations in human motion via mix-and-match perturbation. *arXiv preprint arXiv:1908.00733*, 2019.
- Emad Barsoum, John Kender, and Zicheng Liu. Hp-gan: Probabilistic 3d human motion prediction via gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1418–1427, 2018.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*, 2015.
- Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- Jaemin Cho, Minjoon Seo, and Hannaneh Hajishirzi. Mixture content selection for diverse sequence generation. *arXiv preprint arXiv:1909.01953*, 2019.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Tim R Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M Tomczak. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018.
- Adji B Dieng, Yoon Kim, Alexander M Rush, and David M Blei. Avoiding latent variable collapse with generative skip models. *arXiv preprint arXiv:1807.04863*, 2018.
- Le Fang, Chunyuan Li, Jianfeng Gao, Wen Dong, and Changyou Chen. Implicit deep latent variable models for text generation. *arXiv preprint arXiv:1908.11527*, 2019.
- Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4346–4354, 2015.
- Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *2017 International Conference on 3D Vision (3DV)*, pp. 458–466. IEEE, 2017.
- Anirudh Goyal ALIAS PARTH Goyal, Alessandro Sordoni, Marc-Alexandre Côté, Nan Rosemary Ke, and Yoshua Bengio. Z-forcing: Training stochastic recurrent networks. In *Advances in neural information processing systems*, pp. 6713–6723, 2017.
- Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 786–803, 2018a.
- Liang-Yan Gui, Yu-Xiong Wang, Deva Ramanan, and José MF Moura. Few-shot human motion prediction via meta-learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 432–450, 2018b.
- Ishaan Gulrajani, Kundan Kumar, Faruk Ahmed, Adrien Ali Taiga, Francesco Visin, David Vazquez, and Aaron Courville. Pixelvae: A latent variable model for natural images. *arXiv preprint arXiv:1611.05013*, 2016.

- Kelvin Guu, Tatsunori B Hashimoto, Yonatan Oren, and Percy Liang. Generating sentences by editing prototypes. *Transactions of the Association of Computational Linguistics*, 6:437–450, 2018.
- Junxian He, Daniel Spokoyny, Graham Neubig, and Taylor Berg-Kirkpatrick. Lagging inference networks and posterior collapse in variational autoencoders. *arXiv preprint arXiv:1901.05534*, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014.
- Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5308–5317, 2016.
- Yoon Kim, Sam Wiseman, Andrew C Miller, David Sontag, and Alexander M Rush. Semi-amortized variational autoencoders. *arXiv preprint arXiv:1802.02550*, 2018.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Manoj Kumar, Mohammad Babaeizadeh, Dumitru Erhan, Chelsea Finn, Sergey Levine, Laurent Dinh, and Durk Kingma. Videoflow: A flow-based generative model for video. *arXiv preprint arXiv:1903.01434*, 2019.
- Jogendra Nath Kundu, Maharshi Gor, and R Venkatesh Babu. Bihmp-gan: Bidirectional 3d human motion prediction gan. *arXiv preprint arXiv:1812.02591*, 2018.
- Bohan Li, Junxian He, Graham Neubig, Taylor Berg-Kirkpatrick, and Yiming Yang. A surprisingly effective fix for deep latent variable modeling of text. *arXiv preprint arXiv:1909.00868*, 2019.
- Chao Li, Qiaoyong Zhong, Di Xie, and Shiliang Pu. Co-occurrence feature learning from skeleton data for action recognition and detection with hierarchical aggregation. *arXiv preprint arXiv:1804.06055*, 2018.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Xiao Lin and Mohamed R Amer. Human motion modeling using dvgans. *arXiv preprint arXiv:1804.10652*, 2018.
- Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, Lawrence Carin, et al. Cyclical annealing schedule: A simple approach to mitigating kl vanishing. *arXiv preprint arXiv:1903.10145*, 2019.
- Thomas Lucas and Jakob Verbeek. Auxiliary guided autoregressive variational autoencoders. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 443–458. Springer, 2018.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

- Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *ICCV*, 2019.
- Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4674–4683. IEEE, 2017.
- Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. *arXiv preprint arXiv:1805.06485*, 2018.
- Dario Pavllo, Christoph Feichtenhofer, Michael Auli, and David Grangier. Modeling human motion with quaternion-based neural networks. *arXiv preprint arXiv:1901.07677*, 2019.
- Ali Razavi, Aäron van den Oord, Ben Poole, and Oriol Vinyals. Preventing posterior collapse with delta-vaes. *arXiv preprint arXiv:1901.03416*, 2019.
- Cristian Rodriguez, Basura Fernando, and Hongdong Li. Action anticipation by predicting future dynamic images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 0–0, 2018.
- Mohammad Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Basura Fernando, Lars Petersson, and Lars Andersson. Encouraging lstms to anticipate actions very early. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- Stanislau Semeniuta, Aliaksei Severyn, and Erhardt Barth. A hybrid convolutional variational autoencoder for text generation. *arXiv preprint arXiv:1702.02390*, 2017.
- Tianxiao Shen, Myle Ott, Michael Auli, and Marc’Aurelio Ranzato. Mixture models for diverse machine translation: Tricks of the trade. *arXiv preprint arXiv:1902.07816*, 2019.
- Yuge Shi, Basura Fernando, and Richard Hartley. Action anticipation with rbf kernelized feature mapping rnn. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 301–317, 2018.
- Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein autoencoders. *arXiv preprint arXiv:1711.01558*, 2017.
- Aaron van den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017.
- Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pp. 3352–3361. IEEE, 2017.
- Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- Jiacheng Xu and Greg Durrett. Spherical latent spaces for stable variational autoencoders. *arXiv preprint arXiv:1808.10805*, 2018.
- Xinchen Yan, Akash Rastogi, Ruben Villegas, Kalyan Sunkavalli, Eli Shechtman, Sunil Hadap, Ersin Yumer, and Honglak Lee. Mt-vae: Learning motion transformations to generate multimodal human dynamics. In *European Conference on Computer Vision*, pp. 276–293. Springer, 2018.
- Dingdong Yang, Seunghoon Hong, Yunseok Jang, Tiangchen Zhao, and Honglak Lee. Diversity-sensitive conditional generative adversarial networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJliMh09F7>.
- Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, and Taylor Berg-Kirkpatrick. Improved variational autoencoders for text modeling using dilated convolutions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 3881–3890. JMLR. org, 2017.
- Ye Yuan and Kris Kitani. Diverse trajectory forecasting with determinantal point processes. *arXiv preprint arXiv:1907.04967*, 2019.

Weiyu Zhang, Menglong Zhu, and Konstantinos G Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2248–2255, 2013.

Shengjia Zhao, Jiaming Song, and Stefano Ermon. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017.

A DETAILED TECHNICAL BACKGROUND ON EVIDENCE LOWER BOUND

To solve the maximum likelihood problem, we would like to have $p_\theta(x|z)$ and $p_\theta(z)$. Using *Variational Inference*, we aim to approximate the true posterior $p_\theta(z|x)$ with another distribution $q_\phi(z|x)$. This distribution is computed via another neural network parameterized by ϕ (called variational parameters), such that $q_\phi(z|x) \simeq p_\theta(z|x)$. Using such an approximation, *Variational Autoencoders* (Kingma & Welling, 2013), or VAEs in short, are able to optimize the marginal likelihood in a tractable way. The optimization objective of the VAEs is a variational lower bound, also known as evidence lower bound, or ELBO in short. Recall that variational inference aims to find an approximation of the posterior that represents the true one. One way to do this is to minimize the divergence between the approximate and the true posterior using Kullback-Leibler divergence, or KL divergence in short. That is,

$$\mathcal{D}_{KL} [q_\phi(z|x) || p_\theta(z|x)] = \sum_{z \sim q_\phi(z|x)} q_\phi(z|x) \log \frac{q_\phi(z|x)}{p_\theta(z|x)} \quad (9)$$

This can be seen as an expectation,

$$\begin{aligned} \mathcal{D}_{KL} [q_\phi(z|x) || p_\theta(z|x)] &= \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log \frac{q_\phi(z|x)}{p_\theta(z|x)} \right] \\ &= \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log q_\phi(z|x) - \log p_\theta(z|x) \right] \end{aligned} \quad (10)$$

The second term above, i.e., the true posterior, can according to Bayes' theorem, be written as $p_\theta(z|x) = \frac{p_\theta(x|z)p(z)}{p_\theta(x)}$. The data distribution $p_\theta(x)$ is independent of the latent variable z , and can thus be pulled out of the expectation term,

$$\mathcal{D}_{KL} [q_\phi(z|x) || p_\theta(z|x)] = \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log q_\phi(z|x) - \log p_\theta(x|z) - \log p(z) \right] + \log p_\theta(x) \quad (11)$$

By shifting the $\log p_\theta(x)$ term to the right hand side of the above equation, we can write,

$$\begin{aligned} \mathcal{D}_{KL} [q_\phi(z|x) || p_\theta(z|x)] - \log p_\theta(x) &= \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log q_\phi(z|x) - \log p_\theta(x|z) - \log p(z) \right] \\ \log p_\theta(x) - \mathcal{D}_{KL} [q_\phi(z|x) || p_\theta(z|x)] &= \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log p_\theta(x|z) - (\log q_\phi(z|x) - \log p(z)) \right] \\ &= \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log p_\theta(x|z) \right] - \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log q_\phi(z|x) - \log p(z) \right] \end{aligned} \quad (12)$$

The second expectation term in the above equation is, by definition, the KL divergence between the approximate posterior $q_\phi(z|x)$ and the prior $\log p(z)$ distributions. Thus, this can be written as

$$\log p_\theta(x) - \mathcal{D}_{KL} [q_\phi(z|x) || p_\theta(z|x)] = \mathbb{E}_{z \sim q_\phi(z|x)} \left[\log p_\theta(x|z) \right] - \mathcal{D}_{KL} [q_\phi(z|x) || p(z)] \quad (13)$$

In the above equation, $\log p_\theta(x)$ is the log-likelihood of the data which we would like to optimize. $D_{KL}[q_\phi(z|x)||p_\theta(z|x)]$ is the KL divergence between the approximate and the true posterior distributions, and while not computable, from definition we know that it is non-negative. $\mathbb{E}_{z \sim q_\phi(z|x)}[\log p_\theta(x|z)]$ is the reconstruction loss, and $D_{KL}[q_\phi(z|x)||p(z)]$ is the KL divergence between the approximate posterior distribution and a prior over the latent variable. The last term can be seen as a regularizer of the latent representation. Therefore, the intractability and non-negativity of $D_{KL}[q_\phi(z|x)||p_\theta(z|x)]$ only allows us to optimize the lower bound of the log-likelihood of the data,

$$\log p_\theta(x) \geq \mathbb{E}_{z \sim q_\phi(z|x)}[\log p_\theta(x|z)] - \mathcal{D}_{KL}[q_\phi(z|x)||p(z)] \quad (14)$$

which we call variational or evidence lower bound (ELBO).

B KL DIVERGENCE BETWEEN TWO GAUSSIAN DISTRIBUTIONS

In our approach, the model encourages the posterior of CPP-VAE to be close to the one of the CS-VAE. In general, the KL divergence between two distributions P_1 and P_2 is defined as

$$\mathcal{D}_{KL}(P_1||P_2) = \mathbb{E}_{P_1} \left[\log \frac{P_1}{P_2} \right] \quad (15)$$

In a general case, one can have a multivariate Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ in \mathbb{R}^d where $\Sigma = \text{diag}(\sigma^2)$ where σ and μ are predicted by the encoder network of the VAE. The density function of such a distribution is

$$p(x) = \frac{1}{(2\pi)^{\frac{d}{2}} \det(\Sigma)^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right) \quad (16)$$

Thus, the KL divergence between two multivariate Gaussians is computed as

$$\begin{aligned} \mathcal{D}_{KL}(P_1||P_2) &= \frac{1}{2} \mathbb{E}_{P_1} \left[-\log \det \Sigma_1 - (x - \mu_1)^T \Sigma_1^{-1}(x - \mu_1) + \log \det \Sigma_2 + (x - \mu_2)^T \Sigma_2^{-1}(x - \mu_2) \right] \\ &= \frac{1}{2} \log \frac{\det \Sigma_2}{\det \Sigma_1} + \frac{1}{2} \mathbb{E}_{P_1} \left[-(x - \mu_1)^T \Sigma_1^{-1}(x - \mu_1) + (x - \mu_2)^T \Sigma_2^{-1}(x - \mu_2) \right] \\ &= \frac{1}{2} \log \frac{\det \Sigma_2}{\det \Sigma_1} + \frac{1}{2} \mathbb{E}_{P_1} \left[-\text{tr}\{\Sigma_1^{-1}(x - \mu_1)(x - \mu_1)^T\} + \text{tr}\{\Sigma_2^{-1}(x - \mu_2)(x - \mu_2)^T\} \right] \\ &= \frac{1}{2} \log \frac{\det \Sigma_2}{\det \Sigma_1} + \frac{1}{2} \mathbb{E}_{P_1} \left[-\text{tr}\{\Sigma_1^{-1} \Sigma_1\} + \text{tr}\{\Sigma_2^{-1}(xx^T - 2x\mu_2^T + \mu_2\mu_2^T)\} \right] \\ &= \frac{1}{2} \log \frac{\det \Sigma_2}{\det \Sigma_1} - \frac{1}{2}d + \frac{1}{2} \text{tr}\{\Sigma_2^{-1}(\Sigma_1 + \mu_1\mu_1^T - 2\mu_2\mu_1^T + \mu_2\mu_2^T)\} \\ &= \frac{1}{2} \left[\log \frac{\det \Sigma_2}{\det \Sigma_1} - d + \text{tr}\{\Sigma_2^{-1} \Sigma_1\} + \text{tr}\{\mu_1^T \Sigma_2^{-1} \mu_1 - 2\mu_1^T \Sigma_2^{-1} \mu_2 + \mu_2^T \Sigma_2^{-1} \mu_2\} \right] \\ &= \frac{1}{2} \left[\log \frac{|\Sigma_2|}{|\Sigma_1|} - d + \text{tr}\{\Sigma_2^{-1} \Sigma_1\} + (\mu_2 - \mu_1)^T \Sigma_2^{-1} (\mu_2 - \mu_1) \right]. \end{aligned} \quad (17)$$

where $\text{tr}\{\cdot\}$ is the trace operation. In Eq. 17, the covariance matrix Σ_1 and mean μ_1 corresponds to distribution P_1 and covariance matrix Σ_2 and mean μ_2 corresponds to distribution P_2 .

$$\mathcal{L}_{KL} = -\frac{1}{2} \left[\log \frac{|\Sigma_2|}{|\Sigma_1|} - d + \text{tr}\{\Sigma_2^{-1} \Sigma_1\} + (\mu_2 - \mu_1)^T \Sigma_2^{-1} (\mu_2 - \mu_1) \right]. \quad (18)$$

Given Eq. 18, we can then compute the KL divergence of the CPP-VAE and the posterior distribution with mean $\mu + \sigma \odot \mu_c$ and covariance matrix $\text{diag}((\sigma \odot \sigma_c)^2)$. Let $\Sigma = \text{diag}(\sigma^2)$, $\Sigma_c = \text{diag}(\sigma_c^2)$,

d the dimensionality of the latent space, and $tr\{\cdot\}$ the trace of a square matrix. The loss in Eq. (6) can then be written as⁴

$$\mathcal{L}_{prior}^{CPP-VAE} = -\frac{1}{2} \left[\log \frac{|\Sigma_c|}{|\Sigma_c||\Sigma|} - d + tr\{\Sigma_c^{-1}\Sigma_c\Sigma\} + (\mu_c - (\mu + \Sigma\mu_c))^T \Sigma_c^{-1} (\mu_c - (\mu + \Sigma\mu_c)) \right]. \quad (19)$$

Since $\Sigma_c^{-1}\Sigma_c = I$, $|\Sigma_c|$ will be cancelled out in the log term, which yields

$$\mathcal{L}_{prior}^{CPP-VAE} = -\frac{1}{2} \left[\log \frac{1}{|\Sigma|} - d + tr\{\Sigma\} + (\mu_c - (\mu + \Sigma\mu_c))^T \Sigma_c^{-1} (\mu_c - (\mu + \Sigma\mu_c)) \right]. \quad (20)$$

C STOCHASTIC HUMAN MOTION PREDICTION RELATED WORK

Most motion prediction methods are based on *deterministic* models (Pavlo et al., 2018; 2019; Gui et al., 2018a; Jain et al., 2016; Martinez et al., 2017; Gui et al., 2018b; Fragkiadaki et al., 2015; Ghosh et al., 2017), casting motion prediction as a regression task where only one outcome is possible given the observation. While this may produce accurate predictions, it fails to reflect the stochastic nature of human motion, where multiple futures can be highly likely for a single given series of observations. Modeling stochasticity is the topic of this paper, and we therefore focus the discussion below on the other methods that have attempted to do so.

The general trend to incorporate variations in the predicted motions consists of combining information about the observed pose sequence with a random vector. In this context, two types of approaches have been studied: The techniques that directly incorporate the random vector into the RNN decoder and those that make use of an additional CVAE. In the first class of methods, (Lin & Amer, 2018) samples a random vector $z_t \sim \mathcal{N}(0, I)$ at each time step and adds it to the pose input of the RNN decoder. By relying on different random vectors at each time step, however, this strategy is prone to generating discontinuous motions. To overcome this, (Kundu et al., 2018) makes use of a single random vector to generate the entire sequence. This vector is both employed to alter the initialization of the decoder and concatenated with a pose embedding at each iteration of the RNN. By relying on concatenation, these two methods contain parameters that are specific to the random vector, and thus give the model the flexibility to ignore this information. In (Barsoum et al., 2018), instead of using concatenation, the random vector is added to the hidden state produced by the RNN encoder. While addition prevents having parameters that are specific to the random vector, this vector is first transformed by multiplication with a learnable parameter matrix, and thus can again be zeroed out so as to remove the source of diversity, as observed in our experiments. The second category of stochastic methods introduce an additional CVAE between the RNN encoder and decoder. This allows them to learn a more meaningful transformation of the noise, combined with the conditioning variables, before passing the resulting information to the RNN decoder. In this context, (Walker et al., 2017) proposes to directly use the pose as conditioning variable. As will be shown in our experiments, while this approach is able to maintain some degree of diversity, albeit less than ours, it yields motions of lower quality because of its use of independent random vectors at each time step. Instead of perturbing the pose, the recent work (Yan et al., 2018) uses the RNN decoder hidden state as conditioning variable in the CVAE, concatenating it with the random vector. While this approach generates high-quality motions, it suffers from the fact that the CVAE decoder gives the model the flexibility to ignore the random vector, which therefore yields low-diversity outputs. Similar to (Yan et al., 2018) Mix-and-Match (Aliakbarian et al., 2019) perturbs the hidden states, but replaces the deterministic concatenation operation with a stochastic perturbation of the hidden state with the noise. Through such a perturbation, the decoder is not able to decouple the noise and the condition, the phenomenon that happens in concatenation (Yan et al., 2018). However, since the perturbation is not learned and is a non-parametric operation, the quality of generated motion is comparably low.

Generating diverse plausible motions given limited observations has many applications, especially when the motions are generated in an action-agnostic manner, as done in our work. For instance, our model can be used for human action forecasting (Rodriguez et al., 2018; Sadegh Aliakbarian et al., 2017; Shi et al., 2018; Aliakbarian et al., 2018), where one seeks to anticipate the action as early as possible, where one modality utilized is human motion/poses.

⁴See Appendix B for more details on the KL divergence between two multivariate Gaussians.

D FURTHER DISCUSSION ON THE PERFORMANCE OF STOCHASTIC BASELINES

The MT-VAE model (Yan et al., 2018) tends to ignore the random variable z , thus ignoring the root of variation. As a consequence, it achieves a low diversity, much lower than ours, but produces samples of high quality, albeit almost identical (see the qualitative comparison of different baselines in the appendix). To further confirm that the MT-VAE ignores the latent variable, we performed an additional experiment where, at test time, we sampled each element of the random vector independently from $\mathcal{N}(50, 50)$ instead of from the prior $\mathcal{N}(0, I)$. This led to neither loss of quality nor increase of diversity of the generated motions. Experiments on HP-GAN model (Barsoum et al., 2018) evidences the limited diversity of the sampled motions despite its use of random noise during inference. Note that the authors of (Barsoum et al., 2018) mentioned in their paper that the random noise was added to the hidden state. Only by studying their publicly available code⁵ did we understand the precise way this combination was done. In fact, the addition relies on a parametric, linear transformation of the noise vector. That is, the perturbed hidden state is obtained as

$$h_{\text{perturbed}} = h_{\text{original}} + W^{z \rightarrow h} z . \quad (21)$$

Because the parameters $W^{z \rightarrow h}$ are *learned*, the model has the flexibility to ignore z , which leads to the low diversity of sampled motions. Note that the authors of (Barsoum et al., 2018) acknowledged that, despite their best efforts, they noticed very little variation between predictions obtained with different z values. Since the perturbation is ignored, however, the quality of the generated motions is high. The other baseline, Pose-Knows (Walker et al., 2017), produces motions with higher diversity than the aforementioned two baselines, but of much lower quality. The main reason behind this is that the random vectors that are concatenated to the poses at each time-step are sampled independently of each other, which translates to discontinuities in the generated motions. This problem might be mitigated by sampling the noise in a time-dependent, autoregressive manner, as in (Kumar et al., 2019) for video generation. Doing so, however, goes beyond the scope of our analysis. The Mix-and-Match approach (Aliakbarian et al., 2019) yields sampled motions with higher diversity and reasonable quality. The architecture of Mix-and-Match is very close to that of MT-VAE, but replaces the deterministic concatenation operation with a stochastic perturbation of the hidden state with the noise. Through such a perturbation, the decoder is not able decouple the noise and the condition, the phenomenon that happens in concatenation. However, since the perturbation is not learned and is a non-parametric operation, the quality of the generated motion is lower than ours and of other baselines (except for Pose-Knows). We see Mix-and-Match perturbation as a workaround to the posterior collapse problem while sacrificing the quality and the context in the sampled motions. We also provide a more complete related work on diverse human motion prediction in Appendix C.

E ABLATION STUDY ON DIFFERENT MEANS OF CONDITIONING

In addition to the experiments in the main paper, we also study various designs to condition the VAE encoder and decoder. As discussed before, conditioning the VAE encoder can be safely done via concatenating two deterministic sources of information, i.e., the representations of the past and the future, since both sources are useful to compress the future motion into the latent space. In Table 6, we use both a deterministic representation of the observation, h_t , and a stochastic one, z_c , as a conditioning variable for the encoder. Similarly, we compare the use of either of these variables via concatenation with that of our modified reparameterization trick (Eq. 4). This shows that, to condition the decoder, reparameterization is highly effective at addressing posterior collapse. Furthermore, for the encoder, a deterministic condition works better than a stochastic one. When both the encoder and decoder are conditioned via deterministic conditioning variables, i.e., row 2 in Table 6, the model learns to ignore the latent variable and rely solely on the condition, as evidenced by the KL term tending to zero.

F EXPERIMENTAL RESULTS ON PENN ACTION DATASET

As a complementary experiment, we evaluate our approach on the Penn Action dataset, which contains 2326 sequences of 15 different actions, where for each person, 13 joints are annotated in 2D

⁵<https://github.com/ebarsoum/hpgan>

Table 6: Evaluation of various architecture designs for a CVAE. A smaller KL value, indicating posterior collapse, leads to less diversity.

Encoder Conditioning	Decoder Conditioning	CPP-VAE’s Training KL
Concatenation (z_c)	Reparameterization (z_c)	6.92
Concatenation (h_t)	Concatenation (h_t)	0.04
Concatenation (z_c)	Concatenation (z_c)	0.61
Concatenation (h_t)	Reparameterization (z_c)	8.07

space. Most sequences have less than 50 frames and the task is to generate the next 35 frames given the first 15. Results are provided in Table 7. Note that the upper bound for the Context metric is 0.74, i.e., the classification performance given the Penn Action ground-truth motions.

Table 7: Quantitative evaluation on the Penn Action dataset. Note, the diversity of 1.21 is reasonably high for normalized 2D joint positions, i.e., values between 0 and 1, normalized with the width and the height of the image.

Method	ELBO (KL) (Reconstructed)	Diversity (Sampled)	Quality (Sampled)	Context (Sampled)	Training KL (Reconstructed)
CPP-VAE	0.034 (6.07)	1.21	0.46	0.70	4.84
Autoregressive Counterpart	0.048 (N/A)	0.00	0.46	0.51	N/A

G PSEUDO-CODE FOR CPP-VAE

Here, we provide the forward pass pseudo-codes for both CS-VAE and CPP-VAE.

Algorithm 1 A forward pass of CS-VAE

1: procedure CS-VAE(<i>condition</i>)	▷ Human motion up to time t or the source text
2: $h_t = \text{EncodeCondition}(x_t)$	▷ Observed motion/source text encoder
3: $\mu_c, \sigma_c = \text{CS-VAE.Encode}(h_t)$	
4: Sample $\epsilon \sim \mathcal{N}(0, I)$	▷ Sample from standard Gaussian
5: $z_c = \mu_c + \sigma_c \odot \epsilon$	▷ Reparameterization
6: $\hat{h}_t = \text{CS-VAE.Decode}(z_c)$	
7: $\hat{x}_t = \text{DecodeCondition}(\hat{h}_t, \text{seed})$	
8: return $\hat{x}_t, \mu_c, \sigma_c, h_t, z_c$	▷ h_t and z_c condition CPP-VAE encoder and decoder respectively

Algorithm 2 A forward pass of CPP-VAE

1: procedure CPP-VAE(x_T, z_c, h_t)	▷ Human motion from t to T or the target text
2: if isTraining then	
3: $h_T = \text{EncodeData}(x_T)$	▷ Future motion/target sentence encoder
4: $h_{Tt} = \text{Concatenate}(h_T, h_t)$	
5: $\mu, \sigma = \text{CPP-VAE.Encode}(h_{Tt})$	
6: $z = \mu + \sigma \odot z_c$	▷ Our extended reparameterization
7: else	
8: $z = z_c$	
9: $\hat{h}_T = \text{CPP-VAE.Decode}(z)$	
10: $\hat{x}_T = \text{DecodeData}(\hat{h}_T, \text{seed})$	
11: return \hat{x}_T, μ, σ	

H STOCHASTIC HUMAN MOTION PREDICTION ARCHITECTURE

Our motion prediction model follows the architecture depicted in Fig. 1 (a). Below, we describe the architecture of each component in our model. Note that human poses, consisting of 32 joints in case of the Human3.6M dataset, are represented in 4D quaternion space. Thus, each pose at each

time-step is represented with a vector of size 1×128 . All the tensor sizes described below ignores the mini-batch dimension for simplicity.

Observed motion encoder, or the CS-VAE’s motion encoder, is a single layer GRU (Chung et al., 2014) network with 1024 hidden units. If the observation sequence has the length T_{obs} , the observed motion encoder maps $T_{obs} \times 128$ into a single hidden representation of size 1×1024 , i.e., the hidden state of the last time-step. This hidden state, h_t , acts as the condition to the CPP-VAE’s encoder and the direct input to the CS-VAE’s encoder.

CS-VAE, similar to any variational autoencoder, has an encoder and decoder. The CS-VAE’s encoder is a fully-connected network with ReLU non-linearities, mapping the hidden state of the motion encoder, i.e., h_t , to an embedding of size 1×512 . Then, to generate the mean and standard deviation vectors, two fully connected branches are considered. These map the embedding of size 1×512 to a vector of means of size 1×128 and a vector of standard deviation of size 1×128 , where 128 is the length of the latent variable. Note that we apply a ReLU non-linearity to the vector of standard deviations to make sure it is non-negative. We then use the reparameterization trick (Kingma & Welling, 2013) to sample a latent variable of size 1×128 . The CS-VAE’s decoder consists of multiple fully-connected layers, mapping the latent variable to a variable of size 1×1024 , acting as the initial hidden state of the observed motion decoder. Note that, we apply a Tanh non-linearity to the generated hidden state to mimic the properties of a GRU hidden state.

Observed motion decoder, or the CS-VAE’s motion decoder, is similar to its motion encoder, except for the fact that it reconstructs the motion autoregressively. Additionally, it is initialized with the reconstructed hidden state, i.e., the output of CS-VAE’s decoder. The output of each GRU cell at each time-step is then fed to a fully-connected layer, mapping the GRU output to a vector of size 1×128 which represents a human pose with 32 joints in 4D quaternion space. To decode the motions, we use a teacher forcing technique (Williams & Zipser, 1989) during training. At each time-step, the network chooses with probability P_{tf} whether to use its own output at the previous time-step or the ground-truth pose as input. We initialize $P_{tf} = 1$, and decrease it linearly at each training epoch such that, after a certain number of epochs, the model becomes completely autoregressive, i.e., uses only its own output as input to the next time-step. Note, at test time, motions are generated completely autoregressively, i.e., with $P_{tf} = 0$.

Note, the future motion encoder and decoder have exactly the same architectures as the observed motion ones. The only difference is their input, where the future motion is represented by poses from T_{obs} to T_{end} in a sequence. In the following, we describe the architecture of CPP-VAE for motion prediction.

CPP-VAE is a conditional variational encoder. Its encoder’s input is a representation of future motion, i.e., the last hidden state of the future motion encoder called h_T , conditioned on h_t . The conditioning is done by concatenation, thus, the input to the encoder is a representation of size 1×2048 . The CPP-VAE’s encoder, similar to CS-VAE’s encoder, maps its input representation to an embedding of size 1×512 . Then, to generate the mean and standard deviation vectors, two fully connected branches are considered, mapping the embedding of size 1×512 to a vector of means of size 1×128 and a vector of standard deviations of size 1×128 , where 128 is the length of the latent variable. Note that we apply a ReLU non-linearity to the vector of standard deviations to make sure it is non-negative. To sample the latent variable, we use our extended reparameterization trick, explained in Eq. 4. This unifies the conditioning and sampling of the latent variable. Then, similar to CS-VAE, the latent variable is fed to the CPP-VAE’s decoder, which is a fully connected network that maps the latent representation of size 1×128 to a reconstructed hidden state of size 1×1024 for future motion. Note that, we apply a Tanh non-linearity to the generated hidden state to mimic the properties of a GRU hidden state.

I DIVERSE IMAGE CAPTIONING ARCHITECTURE

Our diverse image captioning model follows the architecture depicted in Fig. 1 (a). Below, we describe the architecture of each component in our model. Note, all tensor sizes described below ignore the mini-batch dimension for simplicity.

Image encoder is, here, ResNet152 (He et al., 2016) pretrained on ImageNet (Krizhevsky et al., 2012). Given the encoder, the conditioning signal is a 1×2048 feature representation. Note that,

to avoid an undesirable equilibrium in the reconstruction loss of the CS-VAE, we freeze ResNet152 during training.

CS-VAE is a standard variational autoencoder. The encoder of the CS-VAE maps the input representation of size 1×2048 to an embedded representation of size 1×1024 . Then, to generate the mean and standard deviation vectors, two fully connected branches are considered, mapping the embedding of size 1×1024 to a vector of means of size 1×256 and a vector of standard deviations of size 1×256 , where 256 is the length of the latent variable. The decoder of the CS-VAE maps the sampled latent variable of size 1×256 to a representation of size 1×2048 . The generated representation acts as a reconstructed image representation. During training, we learn the reconstruction by computing the smoothed L_1 loss⁶ between the generated representation and the image feature (of the frozen ResNet152).

Caption encoder is a single layer GRU network with the hidden size of 1024. Each word in the caption is represented through a randomly initialized embedding layer that maps each word to a representation of size 1×1024 . The caption encoder gets a caption as input and generates a hidden representation of size 1×1024 .

CPP-VAE is a conditional variational autoencoder. As the input to its encoder, we first concatenate the image representation of size 1×2048 to the caption representation of size 1×1024 . The encoder then maps this representation to an embedded representation of size 1×1024 . Then, to generate the mean and standard deviation vectors, two fully connected branches are considered, mapping the embedding of size 1×1024 to a vector of means of size 1×256 and a vector of standard deviations of size 1×256 , where 256 is the length of the latent variable. To sample the latent variable, we make use of our extended reparameterization trick, explained in Eq. 4. This unifies the conditioning and sampling of the latent variable. The CPP-VAE’s decoder then maps this latent representation to a vector of size 1×1024 through a few fully-connected layers. We then apply a batch normalization (Ioffe & Szegedy, 2015) on the representation which then acts as the first token to the caption decoder.

Caption decoder is also a single layer GRU network with a hidden size of 1024. Its first token is the representation generated by the CPP-VAE’s decoder, while the rest of tokens are represented by the words in the corresponding caption. To decode the caption, we use a teacher forcing technique during training. At each time-step, the network chooses with probability P_{tf} whether to use its own output at the previous time-step or the ground-truth token as input. We initialize $P_{tf} = 1$, and decrease it linearly at each training epoch such that, after a certain number of epochs, the model becomes completely autoregressive, i.e., uses only its own output as input to the next time-step. Note, at test time, motions are generated completely autoregressively, i.e., with $P_{tf} = 0$.

J DIVERSE TEXT GENERATION RELATED WORK

There are a number of studies which utilize generative models for language modeling. For instance, (Fang et al., 2019) uses VAEs and LSTMs in an unconditional language modeling problem where posterior collapse may occur if the VAE is not trained well. To handle the problem of posterior collapse in language modeling, the authors of (Fang et al., 2019) try to directly match the aggregated posterior to the prior. It is discussed that this can be considered an extension of variational autoencoders with a regularization when maximizing mutual information, addressing the posterior collapse issue. VAEs are also used for language modeling in (Li et al., 2019). It was observed that for language modeling with VAEs it is hard to find a good balance between language modeling and representation learning. To improve the training of VAEs in such scenarios, the authors of (Li et al., 2019) first pretrain the inference network in an autoencoder fashion such that the inference network learns a good representation of the data in a deterministic manner. Then, they train the whole VAE while considering a weight for the KL term during training. However, the second step modifies the way VAEs optimize the variational lower bound. The proposed technique also prevents the model from being trained end-to-end.

Unlike these approaches, our method considers the case of conditional sequence (text) generation where the conditioning signal (the image to be captioned in our case) is strong enough such that the caption generator can rely solely on that.

⁶<https://pytorch.org/docs/stable/nn.html#smoothl1loss>

A recent work (Cho et al., 2019) proposes to separate the diversification from generation when it comes to sequence generation and language modeling. The diversification stage uses a mixture of experts (MoE) to sample different binary masks on the source sequence for diverse content selection. The generation stage uses a standard encoder-decoder model given each selected content from the source sequence. While shown to be effective in generating diverse sequences, it relies heavily on the selection part, where one need to select the information in the source that is more important to generate the target sequence. Thus, the diversity of the generated target sequence depends on the diversity of the selected parts of the source sequence. Similarly, the authors of (Shen et al., 2019) utilize MoE for the task of diverse machine translation. While this task is considered to be diverse text generation and shown to be highly successful in generating diverse translations of each source sentence, it relies on the availability of the a stochastic dataset, i.e., having access to multiple target sequences for each source sentence during training.

While these approaches are successful in generating diverse sentences given the conditioned sequence, unlike our approach that works with deterministic datasets, they assume having access to a stochastic dataset.

K ABLATION STUDY ON DIVERSE IMAGE CAPTIONING

In addition to the experiments in the main paper, in Table 8, we also evaluate our approach, as well as the autoregressive baseline and the CVAE, in terms of BLEU score for BLEU1, BLEU2, BLEU3, and BLEU4 of generated captions at test time. For the autoregressive baseline, the model generates one caption per image, thus, it is straightforward to compute the BLEU scores. For the CVAE, we consider the best BLEU score among all $K = 50$ sampled captions according the the best matching ground-truth caption. For our model, we consider the caption from mode, i.e., the one sampled from $z = \mu_c$. Although the caption sampled from CPP-VAE is not chosen based on the best match with the ground-truth caption (similar to CVAE), it shows promising quality in terms of BLEU scores. For the sake of completeness and fairness, we also provide the results with best of K captions for our approach as well.

Table 8: BLEU scores of different orders for sampled captions from our model as well as the baselines.

Model	BLEU1	BLEU2	BLEU3	BLEU4
Autoregressive (deterministic)	0.46	0.39	0.21	0.16
Conditional VAE (best of K captions)	0.44	0.38	0.20	0.17
CPP-VAE (caption from mode)	0.44	0.37	0.20	0.14
CPP-VAE (best of K captions)	0.45	0.39	0.23	0.18

The results in Table 8 clearly shows the effectiveness of sampling from mode in our approach. In this case, one could simply rely on the mode of the distribution to achieve a reasonably high quality caption.

L HUMAN MOTION PREDICTION QUALITATIVE RESULTS

Here we provide a number of qualitative results on diverse human motion prediction on the Human3.6M dataset. As can be seen in Figures 3 to 8, the motions generated by our approach are diverse and natural, and mostly within the context of the observed motion.

M DIVERSE IMAGE CAPTIONING QUALITATIVE RESULTS

In this section, we provide a number of qualitative examples of captions generated by our approach. Illustrated in Figures 9 to 16, there are five different ground-truth captions per image. However, as mentioned in the paper, during training we only utilize one (i.e., training with a deterministic dataset). While captions generated by our approach are diverse, they all describe the image adequately. Note that it is a feature of our approach to generate a caption from the mode of its distribution, usually achieving a good descriptive caption. This is also evidenced by the quantitative results



Figure 3: Qualitative evaluation of the diversity in human motion. The first row illustrates the ground-truth motion. The first six poses of each row depict the observation (the condition) and the rest are sampled from our model. Each row is a randomly sampled motion (not cherry picked). As can be seen, all sampled motions are natural, with a smooth transition from the observed to the generated ones. The diversity increases as we increase the sequence length.



Figure 4: Additional qualitative evaluation of the diversity in human motion.

in Table 8 where the BLEU scores for the caption from mode is relatively high compared to other baselines. Note that for the conditional VAE, all sampled captions are identical, despite sampling multiple latent variables. Therefore, we provide only one caption for this baseline.



Figure 5: Additional qualitative evaluation of the diversity in human motion.

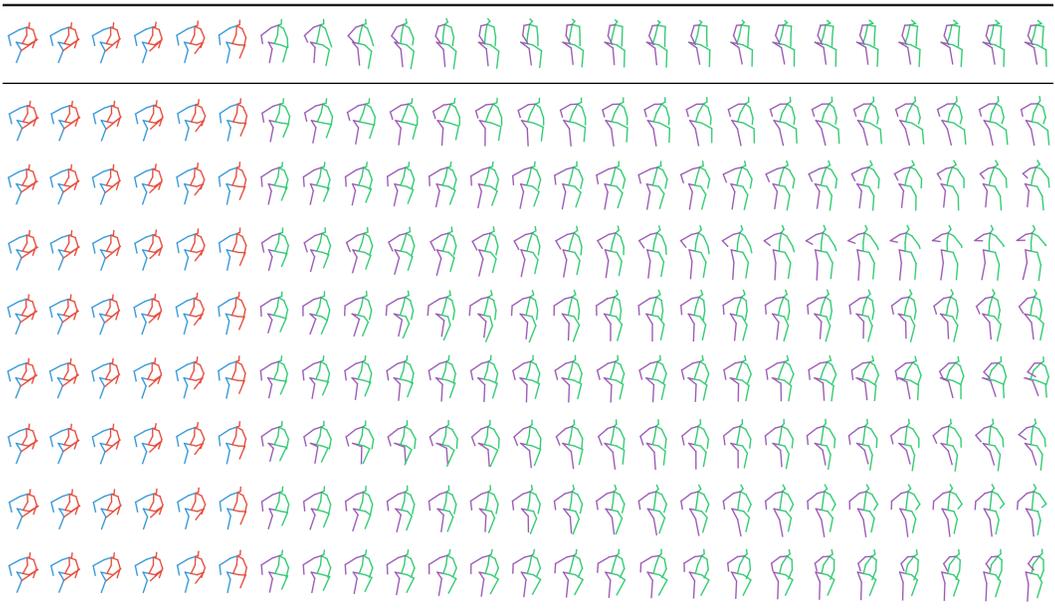


Figure 6: Additional qualitative evaluation of the diversity in human motion.



Figure 7: Additional qualitative evaluation of the diversity in human motion.



Figure 8: Additional qualitative evaluation of the diversity in human motion.



ground-truth Captions:

1. A big blue two level bus at a bus stop with people.
2. a blue and white double decker bus parked at a bus stop
3. A blue bus s parked on a curb where people are standing.
4. A blue double decker bus parked on the side of a road.
5. A bus pulls up to a station where people wait.

Caption from Autoregressive:

a blue bus is parked on the side of the road .

Caption from Conditional VAE:

a bus is parked on the side of the road .

Captions from CPP-VAE's sampling:

1. a large bus is going down the street next to a huge crowd of people .
2. a double decker bus is driving down the street .
3. a blue and white bus on a street next to a building .
4. a bus is driving down the street with a car on the opposite side of the road .
5. a bus is driving down the street with a group of people on the side of it .
6. a street with a lot of people on it and a street sign on it .
7. a group of people standing on a street near a bus .
8. a street sign on a pole with a bus in the background .
9. a bus is parked on the side of the street .

Caption from CPP-VAE's mode:

a double decker bus is driving down the street in pursuit of passengers .

Figure 9: Qualitative evaluation of the diversity in generated captions. While captions generated by our approach are diverse, they all describe the image properly. The caption from mode also usually achieves a good descriptive caption.



ground-truth Captions:

1. A group of sheep standing in a field.
2. Three sheep in a a grassy field with more sheep in the background.
3. A flock of sheep are on a green sloping field.
4. A group of sheep standing in a rocky field of grass.
5. sheep look towards the camera as the stand in a field

Caption from Autoregressive:

a group of sheep standing on top of a lush green field .

Caption from Conditional VAE:

a herd of sheep standing on a lush green hillside .

Captions from CPP-VAE's sampling:

1. a herd of sheep grazing on a lush green hillside .
2. a group of sheep standing around in a field .
3. a herd of sheep standing on top of a grass covered field .
4. a group of sheep are standing in a grassy field .
5. a herd of sheep standing on top of a grass covered hillside .
6. a sheep standing on a hillside with other sheep in the background .
7. a group of sheep grazing in a field .
8. a group of sheep standing on a hill with a view of the grass .
9. a group of sheep standing next to each other .

Caption from CPP-VAE's mode:

a herd of sheep grazing on a lush green hillside .

Figure 10: Additional qualitative evaluation of the diversity in generated captions.



ground-truth Captions:

1. A little girl holds up a big blue umbrella.
2. a young girl stands with her arms wrapped around a large blue umbrella
3. A girl in a pink shirt holding a blue umbrella.
4. A little girl who is holding an umbrella.
5. A little girl with a big, blue umbrella

Caption from Autoregressive:

a woman holding an umbrella standing in the rain .

Caption from Conditional VAE:

a little girl standing in a field holding an umbrella .

Captions from CPP-VAE's sampling:

1. a person holding a rainbow colored umbrella over their head .
2. a woman holding an umbrella in the rain .
3. a woman holding a colorful umbrella over her head .
4. a woman holding an umbrella in the rain .
5. a woman with a red umbrella is standing on the beach .
6. a woman in a red shirt is holding a blue umbrella .
7. a person holding an umbrella while walking down a street .
8. a woman holding an umbrella while walking on a sidewalk .
9. a little girl with a purple umbrella standing in the rain .

Caption from CPP-VAE's mode:

a young woman holding an umbrella standing on a sidewalk .

Figure 11: Additional qualitative evaluation of the diversity in generated captions.



ground-truth Captions:

1. There is a bird that is on a branch sitting looking
2. A small red and gray bird on top of a branch.
3. A bird with an orange beak sitting on a tree branch.
4. A small yellow and red finch resting on a branch.
5. there is a small bird that is standing on the branch

Caption from Autoregressive:

a bird is perched on a tree branch .

Caption from Conditional VAE:

a bird is perched on a branch in a tree .

Captions from CPP-VAE's sampling:

1. a bird sitting on top of a tree branch .
2. a bird is perched on top of a tree branch .
3. a bird is perched on a branch in the sun .
4. a bird is perched on a branch in the shade .
5. a large orange and black bird sitting on top of a metal pole .
6. a bird is standing on the side of a road .
7. a bird is standing on a branch in a tree .
8. a bird is standing on a rock in a patch of grass .
9. a little bird is standing on a branch in a tree .

Caption from CPP-VAE's mode:

a bird sitting on top of a tree branch .

Figure 12: Additional qualitative evaluation of the diversity in generated captions.



ground-truth Captions:

1. A clock and a lamp post on the street with a bridge in the background.
2. A clock sits next to a lamp post on a sidewalk.
3. a street with a lamp post and clock with a bridge behind them
4. An antique pole clock and street lamp adorn a little cobblestone city plaza near an overpass.
5. a pole with a clock on top standing next to a light pole

Caption from Autoregressive:

a clock on a pole near a building .

Caption from Conditional VAE:

a clock on the sidewalk near a city street .

Captions from CPP-VAE's sampling:

1. a clock tower with a clock on it
2. a clock tower with a clock on it 's face .
3. a clock tower with a sky background .
4. a clock on a pole next to a street .
5. a clock on a wall with a plant growing out of it .
6. a clock tower is seen against a blue sky .
7. a clock tower sitting in a garden with a fountain .
8. a clock tower is seen in this picture .
9. a clock tower with a bird on it next to a building .

Caption from CPP-VAE's mode:

a clock tower is reflected in a clear side of the city .

Figure 13: Additional qualitative evaluation of the diversity in generated captions.



ground-truth Captions:

1. Some people are sitting at a table and making sandwiches.
2. A group of people sitting around a table with food.
3. A group of people eating together at a table
4. Group of people sitting down at a table and eating together.
5. A family sits down to a hearty meal.

Caption from Autoregressive:

a group of people sitting around a table eating food .

Caption from Conditional VAE:

a group of people sitting around a table eating .

Captions from CPP-VAE's sampling:

1. a man is holding a knife and fork over a plate of food .
2. a group of people sitting at a table with food .
3. a man and a woman standing next to each other .
4. a little girl sitting at a table with a plate of food .
5. a group of people sitting around a dinner table .
6. a man and a woman eating a donut and drinking a beverage .
7. a man is holding a plate with a sandwich and a knife .
8. a man and a woman are sitting at a table with a plate of food .
9. a group of people sitting around a table with food and beverages .

Caption from CPP-VAE's mode:

a group of people sitting around a table eating food .

Figure 14: Additional qualitative evaluation of the diversity in generated captions.



ground-truth Captions:

1. A woman skiing down a snow covered slope next to two brown dogs.
2. A woman with two dogs skiing cross country.
3. Two yellow dogs are in the snow as a woman is cross country skiing.
4. A woman is cross country skiing along with two dogs in the snow.
5. Two dogs are sitting in the snow with its owner

Caption from Autoregressive:

a man standing in the snow with his skis .

Caption from Conditional VAE:

a man is walking down a snowy hill with a dog .

Captions from CPP-VAE's sampling:

1. a woman on a snow board grinds on a snowy hill .
2. a man in a red jacket and hat is skiing down a hill .
3. a group of dog is laying on a snow covered bank .
4. a man in a red jacket and black pants skiing down a hill
5. a woman on skis in the snow with a dog .
6. a dog with another dog is walking on the snow covered path .
7. a dog is lying down on a snow covered slope .
8. a man is standing in the snow with a snowboard attached to his feet .
9. a dog is walking on the snow with a leash attached to its mouth .

Caption from CPP-VAE's mode:

a couple of dogs are standing on a snowy hill with a person in the background .

Figure 15: Additional qualitative evaluation of the diversity in generated captions.



ground-truth Captions:

1. A hotdog and fork sitting on a plastic plate
2. A hot dog covered in mustard and ketchup.
3. A hot dog with ketchup and mustard is on a styrofoam plate.
4. Cooked snack item in bread on plate with condiment.
5. A hot dog that is laying on a plate.

Caption from Autoregressive:

> a hotdog with toppings on a plate with a pickle .

Caption from Conditional VAE:

> a hotdog with mustard and ketchup on it .

Captions from CPP-VAE's sampling:

1. a hot dog bun filled with mustard and ketchup .
2. a hot dog bun filled with fries and pickles .
3. a hotdog with mustard , ketchup and onion on a bun .
4. a plate of food with hot dog bun filled with cheese and vegetables .
5. a hotdog with mustard , ketchup , mustard , and onions on a bun .
6. a white plate topped with a sandwich and a salad .
7. a plate of food with a hot dog , tomatoes and sausage .
8. a plate with a sandwich and french fries on a table .
9. a plate with a sandwich and french fries .

Caption from CPP-VAE's mode:

a hot dog bun filled with fries and pickles .

Figure 16: Additional qualitative evaluation of the diversity in generated captions.