



HAL
open science

Variational latent discrete representation for time series modelling

Max Cohen, Maurice Charbit, Sylvain Le Corff

► **To cite this version:**

Max Cohen, Maurice Charbit, Sylvain Le Corff. Variational latent discrete representation for time series modelling. IEEE Workshop Statistical Signal Processing, 2023. hal-04140156v3

HAL Id: hal-04140156

<https://hal.science/hal-04140156v3>

Submitted on 3 Jul 2023 (v3), last revised 28 Jul 2023 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - ShareAlike 4.0 International License

Variational Latent Discrete Representation for Time Series Modelling

Max Cohen
Télécom SudParis, CITI, TIPIC
Institut Polytechnique de Paris
maxjcohen@proton.me

Maurice Charbit
Accenta, Boulogne-Billancourt
maurice.charbit@accenta.fr

Sylvain Le Corff
Télécom SudParis, CITI, TIPIC
Institut Polytechnique de Paris

Abstract—Discrete latent space models have recently achieved performance on par with their continuous counterparts in deep variational inference. While they still face various implementation challenges, these models offer the opportunity for a better interpretation of latent spaces, as well as a more direct representation of naturally discrete phenomena. Most recent approaches propose to train separately very high-dimensional prior models on the discrete latent data which is a challenging task on its own. In this paper, we introduce a latent data model where the discrete state is a Markov chain, which allows fast end-to-end training. The performance of our generative model is assessed on a building management dataset and on the publicly available Electricity Transformer Dataset.

Index Terms—Discrete latent space, Recurrent Neural Networks, Variational Inference, ELBO.

I. INTRODUCTION

Discrete latent space models aim at representing data through a finite set of features. Recent advances in generative models have pushed towards these representations, as they fit data with naturally discrete hidden states. For instance, defining a classification task over a dataset often implies a partition of an unobserved latent space, which could be modelled using a categorical random variable, as presented in [1]. In attention-based models, representing the focus location as a discrete variable, i.e. the right place to focus in the past for predicting the next observation, has proven efficient and can help interpreting prediction errors, see for instance [2]. When analyzing time series data, the evolution of a discrete latent variable can be interpreted as a switch in regime, see for instance [3]–[5]. In [6], the authors model the indoor temperature of a building by identifying discrete regimes, such as the opening of a window, the presence of occupants or shade.

As one of the most widespread expression of discrete latent space models, Hidden Markov Models have been successfully applied in various fields ([7]–[9]). Despite modeling hidden states as a discrete Markov chain, they are able to handle complex data structure, see [10], [11] and the references therein for a complete overview. However, dealing with latent data leads to models that are computationally expensive to train, for instance using Expectation Maximization [12] based approaches, and still struggle to handle large scale datasets,

in particular when the models contain high-dimensional additional latent states.

In contrast, deep learning methods are able to infer millions of parameters from huge amounts of data - at the cost of much more complex models - through automatic differentiation and gradient computation. However, discrete variables usually prevent the propagation of the gradient, and in turn straightforward trainings. The Vector-Quantized Variational AutoEncoder (VQ-VAE, [13]) popularized discrete latent spaces for Variational Inference (see [14] for an introduction), by copying the gradient of discrete variables directly to the previous layer. This approximation is detailed in [15]. Additionally, the training of the VQ-VAE is divided into two steps. The first assumes a non informative prior on the latent variables in order to estimate the encoder and decoder parameters. Samples generated after this step cannot catch the complex dependencies that lie in the latent space; to produce coherent samples, a prior model on the latent variables is trained separately, directly on the latent space. Because these prior models are usually already challenging to fit on their own, most applications of the VQ-VAE have not been able to jointly train both parts of the network, see [16]. This paper brings the following contributions.

- We propose a generative model for time series, where the latent space is modelled as a discrete Markov chain, and, conditionally on the latent states, the observations follow a simple autoregressive process. Parameters are jointly estimated by Variational Inference.
- We compared the impact of different prior models to extract high-level features from the input data (based on convolutional and recurrent architectures) on the quality of samples.
- Our model outperforms the state of the art VQ-VAE in both accuracy and computation time on two time series datasets.

II. RELATED WORKS

a) *Discrete latent representation*: one of the most straightforward application of discrete latent representation is derived from semi or unsupervised classification problems, where the data presents distinct semantic classes. The authors of [1] propose to model such data as generated by both a continuous and a categorical class variable. By integrating

a classification mechanism directly in the model, they are able to outperform continuous latent models, while allowing conditional generation. This idea was transposed to Generative Adversarial Networks (introduced in [17]) by the authors of [18], by adding a discrete random variable as the input of the generator. Trainings on MNIST show how the model associates each value with a class of digits, even without supervision.

The interest in quantized latent variables go beyond classification tasks, and appear in attention-based models, which have dominated the Natural Language Processing and Computer Vision fields for the last years. For instance, [2] presents an image captioning model where attention scores parametrize a categorical random variable; by either sampling or computing its expectation, the model extracts visual features used to generate words. Discrete representations have also successfully replaced continuous latent spaces in existing models. In this way, the authors of [16] quantized the latent features of an existing variational recurrent neural network applied to text-to-speech synthesis, and reported improved quality in generated audio samples.

b) Variational Inference: Variational Inference (VI) is a popular approach to estimate such models, which consists in approximating the posterior distribution of the latent states given the observations, within a parametric family of tractable distributions, see [19]–[21]. A trade-off between the expressiveness of the approximated posterior, and its computation cost must be achieved, which represents the main challenge of VI methods. As the computation of the likelihood of the model and its gradient are intractable, the training procedure usually focuses on maximizing a lower bound of the likelihood by gradient ascent. As sampling discrete random variables prevent gradient propagation through the model, a workaround is required to take advantage of automatic differentiation offered by modern deep learning frameworks. The usual alternative to Variational Inference for modelling time dependant data is Sequential Monte Carlo, see for instance [22] and the references therein.

Vector Quantized Variational Auto Encoders introduced a two-stage training for discrete latent models: first, the posterior is approximated through VI while considering an uninformative prior for the discrete latent variables, i.e. a uniform discrete distribution; then, a very expressive prior model is trained on the discrete latent space to produce samples, all other parameters being fixed. Despite its impressive results, the VQ-VAE still suffers from two drawbacks: the differentiation of the posterior and the choice of the prior.

c) Differentiation of the posterior: sampling discrete variables from the posterior prevents gradient propagation through the model; while the authors of the VQ-VAE proposed a straight-through estimator, other approaches have addressed this issue. In [23], the model is optimized through direct loss minimization, a method that introduces additional bias and hyperparameter tuning. Another method, introduced in [24], leverage importance sampling to sample from the posterior, without introducing bias or new parameters. However, a new

differentiable distribution close to the posterior must be introduced, limiting its usage in general settings. The authors of [25] proposed a new differentiable distribution, approximating samples from a categorical law while being differentiable. Although this method does introduce a new hyperparameter, it offers a very appealing trade-off.

d) Prior: during the second stage of the VQ-VAE training, an autoregressive prior is trained on the latent discrete space. In the original paper, the authors chose very resource intensive models for both image [26] and audio [27] datasets. However, more parsimonious prior models already yield encouraging results. For instance, we can find in [16] a comparison of different autoregressive priors on a text-to-speech task, leading to state of the art performances, using a single layered Long Short Term Memory (LSTM) network. On the same task, the authors of [28] were able to perform end-to-end training. Additionally, this prior model was conditioned on text data to produce speech samples.

III. MODEL

Consider a sequence of observations (x_1, \dots, x_T) in \mathbb{R}^d . It is assumed that these observations depend on external observed signals (u_1, \dots, u_T) in \mathbb{R}^q , referred to as commands. In this paper, we also assume that the observations are independent of the commands conditionally on hidden variables (z_1, \dots, z_T) , which take values in a set of codebooks $\mathcal{E} = \{e_1, \dots, e_K\}$. For each of $1 \leq k \leq K$, $e_k \in \mathbb{R}^D$. We then consider the following family of probability density functions:

$$p_\theta(y_{1:T}|u_{1:T}) = \int p_\theta(y_{1:T}|z_{1:T})p_\theta(z_{1:T}|u_{1:T})dz_{1:T}, \quad (1)$$

depending on an unknown parameter $\theta \in \Theta \subset \mathbb{R}^m$ and where for all $1 \leq s \leq t$ and all sequence $\{a_u\}_{u \geq 1}$, $a_{s:t}$ is a shorthand notation for (a_s, \dots, a_t) .

Conditionally on the commands, we assume that the latent states are Markovian and that the conditional law of the observations depends on past latent states, so that our model is more general than Hidden Markov Models. It also differs from other extensions of HMMs with dependencies between the observations, such as the autoregressive processes as described in [29]. In the following paragraphs, we detail the structure of the observation and prior models, as well as the choice of posterior family.

A. Observation model

We consider a Gaussian observation model, and estimate at each time step its mean and variance: $p_\theta(x_{1:T}|z_{1:T}) = \prod_{t=1}^T \varphi_{\mu_t, \sigma_t^2}(x_t)$, where φ_{μ, σ^2} is the probability density function of a Gaussian random variable with mean vector μ and covariance matrix $\sigma^2 I_d$. For all $1 \leq t \leq T$, $\mu_t = g_\theta^\mu(\mu_{t-1}, z_{1:t-1})$ and $\sigma_t = g_\theta^\sigma(\sigma_{t-1}, z_{1:t-1})$, with $\mu_0 \equiv \sigma_0 \equiv 0$. We discuss our choice for the parametric functions g_θ^μ and g_θ^σ in Section V-A.

B. Prior model

We assume that, conditionally on the commands $u_{1:T}$, the latent state is a discrete Markov chain. Write, for all $1 \leq \ell \leq K$, $p_{\theta,1}^\ell = p_\theta(z_1 = e_\ell | u_{1:T})$, and for all $2 \leq t \leq T$ and $1 \leq k, j \leq K$, $p_{\theta,t|t-1}^{k,j} = p_\theta(z_t = e_k | z_{t-1} = e_j, u_{1:T})$. The prior distribution is then defined as:

$$\begin{aligned} \log p_\theta(z_{1:T} | u_{1:T}) &= \sum_{k=1}^K \mathbb{1}_{z_1=e_k} \log p_{\theta,1}^k \\ &+ \sum_{t=2}^T \sum_{j,k=1}^K \mathbb{1}_{z_{t-1}=e_j} \mathbb{1}_{z_t=e_k} \log p_{\theta,t|t-1}^{k,j}. \end{aligned}$$

C. Posterior distribution

As the posterior distribution of $z_{1:T}$ given $x_{1:T}$ is intractable, we use a variational approach to estimate θ . In the original approach proposed in [13], the authors use an encoding function f_φ , depending on an unknown parameter $\varphi \in \mathbb{R}^p$, mapping the observations to a series of encoded latent variables $z_{1:T}^e = f_\varphi(x_{1:T})$. Then, the posterior is approximated by $q_\varphi(z_t = e_k | x_{1:T}) = \mathbb{1}_{e_* = e_k}$, where $e_* = \operatorname{argmin}_{e_\ell \in \mathcal{E}} \|z_t^e - e_\ell\|_2$. We propose, for $1 \leq k \leq K$:

$$q_{\varphi,t}^k = q_\varphi(z_t = e_k | x_{1:T}) \propto \exp\{-\|z_t^e - e_k\|_2^2\}. \quad (2)$$

During both training and inference, this discrete distribution encourages the exploration of the entire set of codebooks. Under the variational distribution, the latent data are assumed to be independent conditionally on the observations.

IV. INFERENCE PROCEDURE

The Evidence Lower BOund (ELBO) can be decomposed in three terms: for all (θ, φ) ,

$$\begin{aligned} \mathcal{L}(\theta, \varphi) &= \mathbb{E}_{q_\varphi} [\log p_\theta(x_{1:T} | z_{1:T})] + \mathbb{E}_{q_\varphi} [\log p_\theta(z_{1:T} | u_{1:T})] \\ &- \mathbb{E}_{q_\varphi} [\log q_\varphi(z_{1:T} | x_{1:T})]. \quad (3) \end{aligned}$$

The last term of the ELBO can be computed explicitly as follows:

$$\mathbb{E}_{q_\varphi} [\log q_\varphi(z_{1:T} | x_{1:T})] = \sum_{t=1}^T \sum_{k=1}^K q_{\varphi,t}^k \log q_{\varphi,t}^k.$$

We approximate the two other terms by drawing $M > 0$ samples under q_φ and computing a Monte Carlo estimator. In order for this operation to be differentiable, we use a reparametrization designed for categorical variables based on the Gumbel-Softmax distribution.

In [30], it is shown that we can draw samples under q_φ by computing $\operatorname{argmax}_{k=1}^K (\log(q_\varphi^k) + g_k)$, where $(g_k)_{k=1}^K$ are independent identically distributed samples from the Gumbel distribution $G(0, 1)$ with probability density $g(x) = e^{-(x+e^{-1})}$. The Gumbel-Softmax distribution [25] aims at using the softmax function as a continuous and differentiable approximation to the argmax operator. We sample (g_1, \dots, g_K) independently from the Gumbel distribution and define, for all $1 \leq k \leq K$, $\pi_{k,t} \propto \exp((\log q_{\varphi,t}^k + g_k)/\tau_t)$, where $\tau_t > 0$

is the softmax temperature, allowing for a smooth interpolation between a Uniform distribution (for large values of τ_t), and a Categorical distribution (for small values of τ_t). We propose to approximate the variational posterior distribution of z_t by the Dirac mass at $\tilde{z}_t = \sum_{k=1}^K \pi_{k,t} e_k$. Through re-parametrization, this method allows for differentiation of the sampled latent vector \tilde{z}_t , with respect to the codebooks e_k , $1 \leq k \leq K$, and the encodings $z_{1:T}^e$. The first term of the ELBO can now be approximated by $(\theta, \phi) \mapsto M^{-1} \sum_{i=1}^M \log p_\theta(\tilde{z}_{1:T} | u_{1:T})$, and the second by $(\theta, \phi) \mapsto M^{-1} \sum_{i=1}^M \log p_\theta(y_{1:T} | \tilde{z}_{1:T})$.

Estimating θ , ϕ and \mathcal{E} jointly can induce instability at the beginning of the training, leading to diminished performances after convergence. In order to control the trade-off between the capacity of the latent state and the observation model, [31] introduced an additional hyperparameter β when computing the ELBO. As shown in [32], we can perform end-to-end training by penalizing the prior and posterior terms by this factor β initialized close to zero, and then slowly increasing its value until reaching $\beta = 1$.

V. EXPERIMENTS

We benchmarked our approach on the public Electricity Transformer Temperature (ETT) Dataset, designed in [33] to forecast Oil temperature based on hourly power load records (ETT_{h1} subset). Models are trained on the first year of available data, and validated on the subsequent four months. Samples from the validation set are presented in Figure 1, where we compare our proposed model with the benchmarked VQ-VAE. We also propose a benchmark on a building management dataset, where we aim at forecasting hourly indoor humidity based on the indoor temperature and outdoor humidity.

A. Chosen architectures

Our prior model can be decomposed in two sub-networks: an input model is responsible for extracting high level features from the commands, while the autoregressive kernel computes the Markov chain transition probabilities. This disentanglement allows us to keep the kernel simple, while still working with high level features:

$$\begin{aligned} \tilde{u}_{1:T} &= f_\theta^{\text{input_model}}(u_{1:T}), \\ h_t &= \ker_\theta(h_{t-1}, \tilde{u}_{t-\Delta:t}), \forall 1 \leq t \leq T, h_0 \equiv 0. \end{aligned}$$

For the input model, we implemented a 3-layered LSTM, with the same latent dimension as the commands. We then compared several auto regressive architectures.

- A simple RNN cell was used as a benchmark, as they struggle to model long term dependencies.
- A Gated Recurrent Unit (GRU) cell, following results in [16]. These architectures have a more refined memory representation, and have shown encouraging results when applied to building management data, as shown in [34].
- A kernel based on causal convolutions, as proposed in [27], where memory is replaced by an explicit dependency on the last $\Delta = 24$ time steps.

For the encoder f_φ and decoder $g_\theta^{\mu,\sigma}$ parametric functions, we chose 3-layered LSTM networks as well.

Finally, we cross validated the number of codebooks using the Root-Mean-Square Error (RMSE), as described in Section V-B, and settled for $K = 8$. The influence of each codebook on the produced samples is illustrated in Figure 2, where we plot the selected codebook during the generation process at each time step. This representation allows for a deeper understanding of the latent representation of the model, while opening the way for further analysis such as regime switching detection.

We experimented with various schedules increasing β from 0 to 1, and found that while this penalization is necessary to estimate all parameters jointly, the choice of schedule had little influence on the performances. Therefore, we increase β linearly between epochs 1 to 100, and keep $\beta = 1$ for the rest of the training, which amounts to 300 total epochs. In all following simulations, the codebooks dimension is fixed to $D = 32$, and the number of Monte Carlo samples to $M = 1$.

We compared our model to a Gaussian discrete linear Hidden Markov Model (HMM) whose parameters are estimated with the Kalman smoother using the Expectation Maximization (EM, [12]) algorithm. We used the same number of hidden states $K = 8$. We also compared our approach to the original VQ-VAE. The architecture of the model is kept similar, to highlight the following difference of methodology and training.

- The autoencoder is trained while considering an uninformative prior. Then, we estimate the parameters of the prior model while the autoencoder is fixed.
- Samples from the posterior distribution are drawn under a Dirac mass, as shown above. In order to compute the gradient, we use the straight-through estimator.

B. Evaluation

Provided a sequence of commands, each benchmark model produces samples predicting the observations. By averaging them over the time steps, we compute the RMSE: let $\hat{x}_{1:T}^i, 1 \leq i \leq N$ be N independent sequences of predictions of $x_{1:T}$, then $\text{RMSE} = (T^{-1} \sum_{t=1}^T (x_t - N^{-1} \sum_{i=1}^N \hat{x}_t^i)^2)^{1/2}$. Additionally, we report the Mean Absolute Error (MAE) criteria: $\text{MAE} = T^{-1} \sum_{t=1}^T \|x_t - N^{-1} \sum_{i=1}^N \hat{x}_t^i\|$. In the presented experiments, we set $N = 100$. Results on the entire validation set are displayed in Table I.

VI. CONCLUSION

In this paper, we introduced a generative model with discrete latent variables for time series. We illustrated our proposed framework using simple prior architectures, which are able to model dependencies in the latent space and allow joint training of all parameters. However, our methodology can also be adapted for much more complex priors, such as the diffusion bridges models presented in [35].

While choosing a suitable number of codebook remains an open question, discrete latents offer new ways of interpreting the model, for instance through regime segmentation.

Fig. 1. We sample $N = 100$ trajectories under the GRU prior, conditioned on a set of commands from the validation split of the ETT dataset, and plot the associated 95% confidence intervals. As a baseline, we added samples from a VQ-VAE model.





TABLE I

COMPARISON OF RMSE, MAE AND COMPUTATION TIME OF OUR MODEL AGAINST THE BENCHMARKED VQ-VAE AND HMM. THE CHOICE OF PRIOR MODEL IMPACTS PERFORMANCES, YET SIMPLE ARCHITECTURES ARE ABLE TO MODEL THE DEPENDENCIES OF THE OBSERVATIONS. MEAN VALUES OF THE ESTIMATORS OVER THE VALIDATION SAMPLES ARE REPORTED WITH THEIR VARIANCE.

ETT dataset	RMSE	MAE	Computation time
Ours (gru)	0.21 ± 0.10	0.17 ± 0.09	39ms
Ours (rnn)	0.21 ± 0.10	0.17 ± 0.09	39ms
Ours (cnn)	0.20 ± 0.10	0.16 ± 0.09	131ms
VQ-VAE	0.28 ± 0.12	0.24 ± 0.12	39ms
HMM	0.44 ± 0.13	0.32 ± 0.10	994ms
Building management dataset	RMSE	MAE	Computation time
Ours (gru)	0.30 ± 0.19	0.24 ± 0.15	39ms
Ours (rnn)	0.30 ± 0.20	0.25 ± 0.17	39ms
Ours (cnn)	0.31 ± 0.20	0.25 ± 0.16	131ms
VQ-VAE	0.55 ± 0.33	0.51 ± 0.32	39ms
HMM	0.46 ± 0.19	0.40 ± 0.17	994ms

Fig. 2. Codebook usage for a week long sample of the ETT dataset. We sample $N = 100$ trajectories from the prior, and plot the most selected codebook index at each time step. The observed oil temperature has been plotted for comparison. The second codebook seems to be correlated with an increase in oil temperature, while the third one is activated during a decrease. The seventh seems to only appear for low oil temperatures.



REFERENCES

- [1] D. P. Kingma, S. Mohamed, D. Jimenez Rezende, and M. Welling, "Semi-supervised learning with deep generative models," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014.
- [2] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. PMLR, 07–09 Jul 2015, pp. 2048–2057.
- [3] Z. Dzunic and J. Fisher III, "Bayesian switching interaction analysis under uncertainty," in *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, S. Kaski and J. Corander, Eds., vol. 33. PMLR, 22–25 Apr 2014, pp. 220–228.
- [4] F. Lindsten, P. Bunch, S. Särkkä, T. B. Schön, and S. J. Godsill, "Rao-blackwellized particle smoothers for conditionally linear gaussian models," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 2, pp. 353–365, 2015.
- [5] N. M. Nguyen, S. Le Corff, and É. Moulines, "Particle rejuvenation of rao-blackwellized sequential monte carlo smoothers for conditionally linear and gaussian models," *EURASIP Journal on Advances in Signal Processing*, vol. 2017, no. 1, pp. 1–15, 2017.
- [6] B. Ajib, S. Lefteriu, A. Caucheteux, and S. Lecoeuche, "Predicting the air temperature of a building zone by detecting different configurations using a switched system identification technique," *Journal of Building Engineering*, vol. 31, p. 100995, 2020.
- [7] D. Wilks, "Multisite generalization of a daily stochastic precipitation generation model," *Journal of Hydrology*, vol. 210, no. 1, pp. 178–191, 1998.
- [8] M. Gales and S. Young, "The application of hidden Markov models in speech recognition," *Foundations and Trends® in Signal Processing*, vol. 1, no. 3, pp. 195–304, 2008.
- [9] T. A. Patterson, A. Parton, R. Langrock, P. G. Blackwell, L. Thomas, and R. King, "Statistical modelling of individual animal movement: an overview of key methods and a discussion of practical challenges," *AStA Advances in Statistical Analysis*, vol. 101, no. 4, pp. 399–438, Oct 2017.
- [10] O. Cappé, É. Moulines, and T. Rydén, "Inference in hidden markov models," in *Springer series in statistics*, 2005.
- [11] R. Douc, E. Moulines, and D. Stoffer, *Nonlinear time series: Theory, methods and applications with R examples*. CRC press, 2014.
- [12] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society: Series B*, vol. 39, pp. 1–38, 1977.
- [13] A. van den Oord, O. Vinyals, and k. kavukcuoglu, "Neural discrete representation learning," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017.
- [14] M. I. Jordan, Z. Ghahramani, T. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine Learning*, vol. 37, pp. 183–233, 2004.
- [15] Y. Bengio, N. Léonard, and A. C. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *ArXiv*, vol. abs/1308.3432, 2013.
- [16] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, A. Rosenberg, B. Ramabhadran, and Y. Wu, "Generating diverse and natural text-to-speech samples using a quantized fine-grained vae and autoregressive prosody prior," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6699–6703.
- [17] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014.
- [18] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.
- [19] R. Salakhudinov and G. Hinton, "Deep boltzmann machines," in *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, D. van Dyk and M. Welling, Eds., vol. 5. PMLR, 2009, pp. 448–455.
- [20] A. Mnih and D. Rezende, "Variational inference for monte carlo objectives," in *Proceedings of The 33rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. F. Balcan and K. Q. Weinberger, Eds., vol. 48. New York, New York, USA: PMLR, 20–22 Jun 2016, pp. 2188–2196.
- [21] Y. Petetin, Y. Janati, and F. Desbouvries, "Structured variational bayesian inference for gaussian state-space models with regime switching," *IEEE Signal Processing Letters*, vol. 28, pp. 1953–1957, 2021.
- [22] M. Cohen, M. Charbit, and S. Le Corff, "Last layer state space model for representation learning and uncertainty quantification," *2023 IEEE Statistical Signal Processing Workshop (SSP)*, 2023.
- [23] G. Lorberbom, A. Gane, T. Jaakkola, and T. Hazan, "Direct optimization through argmax for discrete variational auto-encoder," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alche-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019.
- [24] A. Bartler, F. Wiewel, L. Mauch, and B. Yang, "Training variational autoencoders with discrete latent variables using importance sampling," in *2019 27th European Signal Processing Conference (EUSIPCO)*, 2019, pp. 1–5.
- [25] E. Jang, S. S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," *ArXiv*, vol. abs/1611.01144, 2017.
- [26] A. van den Oord, N. Kalchbrenner, L. Espeholt, k. kavukcuoglu, O. Vinyals, and A. Graves, "Conditional image generation with pixelcnn decoders," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016.
- [27] A. van den Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. van den Driessche, E. Lockhart, L. Cobo, F. Stimberg, N. Casagrande, D. Grewe, S. Noury, S. Dieleman, E. Elsen, N. Kalchbrenner, H. Zen, A. Graves, H. King, T. Walters, D. Belov, and D. Hassabis, "Parallel wavenet: Fast high-fidelity speech synthesis," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 3918–3926.
- [28] Y. Yasuda, X. Wang, and J. Yamagishid, "End-to-end text-to-speech using latent duration based on vq-vae," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 5694–5698.
- [29] R. Douc, Éric Moulines, and T. Rydén, "Asymptotic properties of the maximum likelihood estimator in autoregressive models with Markov regime," *The Annals of Statistics*, vol. 32, no. 5, pp. 2254 – 2304, 2004.
- [30] K. D. Tocher, "Statistical theory of extreme values and some practical applications : A series of lectures," *Journal of the Royal Statistical Society. Series A (General)*, vol. 118, no. 1, pp. 106–106, 1955.
- [31] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-VAE: Learning basic visual concepts with a constrained variational framework," in *International Conference on Learning Representations*, 2017.
- [32] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, "Zero-shot text-to-image generation," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 8821–8831.
- [33] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, "Informer: Beyond efficient transformer for long sequence time-series forecasting," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, pp. 11 106–11 115, May 2021.
- [34] M. Cohen, S. Le Corff, M. Charbit, A. Champagne, G. Nozière, and M. Preda, "End-to-end deep meta modelling to calibrate and optimize energy consumption and comfort," *Energy and Buildings*, vol. 250, p. 111218, 2021.
- [35] M. Cohen, G. Quispe, S. L. Corff, C. Ollion, and E. Moulines, "Diffusion bridges vector quantized variational autoencoders," in *Proceedings of the 39th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Eds., vol. 162. PMLR, 17–23 Jul 2022, pp. 4141–4156.