

**Mansbridge, Alex; Fierimonte, Roberto; Feige, Ilya; Barber, David**

**Improving latent variable descriptiveness by modelling rather than ad-hoc factors.** (English)

Zbl 07097482

Mach. Learn. 108, No. 8-9, 1601-1611 (2019).

Summary: Powerful generative models, particularly in natural language modelling, are commonly trained by maximizing a variational lower bound on the data log likelihood. These models often suffer from poor use of their latent variable, with ad-hoc annealing factors used to encourage retention of information in the latent variable. We discuss an alternative and general approach to latent variable modelling, based on an objective that encourages a perfect reconstruction by tying a stochastic autoencoder with a variational autoencoder (VAE). This ensures by design that the latent variable captures information about the observations, whilst retaining the ability to generate well. Interestingly, although our model is fundamentally different to a VAE, the lower bound attained is identical to the standard VAE bound but with the addition of a simple pre-factor; thus, providing a formal interpretation of the commonly used, ad-hoc pre-factors in training VAEs.

**MSC:**

68T05 Learning and adaptive systems in artificial intelligence

**Keywords:**

generative modelling; latent variable modelling; variational autoencoders; variational inference; natural language processing

**Software:**

darch; PixelVAE; TopicRNN

**Full Text:** DOI

**References:**

- [1] Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A., Jozefowicz, R., & Bengio, S. (2016) Generating sentences from a continuous space. In *Conference on computational natural language learning*.
- [2] Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., et al. (2017). Variational lossy autoencoder. In *International conference on learning representations*.
- [3] Dieng, A. B., Wang, C., Gao, J., & Paisley J. (2017). TopicRNN: A recurrent neural network with long-range semantic dependency. In *International conference on learning representations*.
- [4] Gulrajani, I., Kumar, K., Ahmed, F., Taiga, A. A., Visin, F., Vazquez, D., & Courville, A. (2017). PixelVAE: A latent variable model for natural images. In *International conference on learning representations*.
- [5] Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., & Lerchner, A. (2017). beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International conference on learning representations*.
- [6] Hinton, GE; Salakhutdinov, RR, Reducing the dimensionality of data with neural networks, *Science*, 313, 504-507, (2006) · Zbl 1226.68083
- [7] Hinton, G. E., & Zemel, R. S. (1994). Autoencoders, minimum description length and helmholtz free energy. In *Advances in neural information processing systems*.
- [8] Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. In *International conference on learning representations*.
- [9] Kiros, R., Salakhutdinov, R., & Zemel, R. (2014). Multimodal neural language models. In *International conference on machine learning*.
- [10] Mansimov, E., Parisotto, E., Ba, J. L., & Salakhutdinov, R. (2016). Generating images from captions with attention. In *International conference on learning representations*.
- [11] Pu, Y., Gan, Z., Heno, R., Yuan, X., Li, C., Stevens, A., & Carin, L. (2016). Variational autoencoder for deep learning of images, labels and captions. In *Advances in neural information processing systems*.

- [12] Rezende, D. J., Mohamed, S., \& Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In \textit{International conference on machine learning}.
- [13] Semeniuta, S., Severyn, A., \& Barth, E. (2017). A hybrid convolutional variational autoencoder for text generation. In \textit{Conference on empirical methods in natural language processing}.
- [14] Shah, H., Zheng, B., \& Barber, D. (2017). Generating sentences using a dynamic canvas. In \textit{Association for the advancement of artificial intelligence}.
- [15] Vedantam, R., Fischer, I., Huang, J., \& Murphy, K. (2018). Generative models of visually grounded imagination. In \textit{International conference on learning representations}.
- [16] Wu, M., \& Goodman, N. (2018). Multimodal generative models for scalable weakly-supervised learning. Preprint arXiv:1802.05335.
- [17] Yang, Z., Hu, Z., Salakhutdinov, R., \& Berg-Kirkpatrick, T. (2017). Improved variational autoencoders for text modeling using dilated convolutions. In \textit{International conference on machine learning}.
- [18] Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., \& Fidler, S. (2015). Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In \textit{International conference on computer vision}.

This reference list is based on information provided by the publisher or from digital mathematics libraries. Its items are heuristically matched to zbMATH identifiers and may contain data conversion errors. It attempts to reflect the references listed in the original paper as accurately as possible without claiming the completeness or perfect precision of the matching.