
Improving Variational Auto-Encoders using Householder Flow

Jakub M. Tomczak, Max Welling
University of Amsterdam
J.M.Tomczak@uva.nl, M.Welling@uva.nl

Abstract

Variational auto-encoders (VAE) are scalable and powerful generative models. However, the choice of the variational posterior determines tractability and flexibility of the VAE. Commonly, latent variables are modeled using the normal distribution with a diagonal covariance matrix. This results in computational efficiency but typically it is not flexible enough to match the true posterior distribution. One fashion of enriching the variational posterior distribution is application of *normalizing flows*, *i.e.*, a series of invertible transformations to latent variables with a simple posterior. In this paper, we follow this line of thinking and propose a *volume-preserving flow* that uses a series of *Householder transformations*. We show empirically on MNIST dataset and histopathology data that the proposed flow allows to obtain more flexible variational posterior and highly competitive results comparing to other normalizing flows.

1 Variational Auto-Encoder

Let \mathbf{x} be a vector of D observable variables, $\mathbf{z} \in \mathbb{R}^M$ a vector of stochastic latent units (variables) and let $p(\mathbf{x}, \mathbf{z})$ be a parametric model of the joint distribution. Given N datapoints $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ we typically aim at maximizing the marginal log-likelihood:

$$\ln p(\mathbf{X}) = \sum_{i=1}^N \ln p(\mathbf{x}_i), \quad (1)$$

with respect to parameters. This task could be troublesome due to intractability of the marginal likelihood, *e.g.*, when the model is parameterized by a neural network (NN). To overcome this issue one can introduce an *inference model* (an *encoder*) $q(\mathbf{z}|\mathbf{x})$ and optimize the variational lower bound:

$$\ln p(\mathbf{x}) \geq \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\ln p(\mathbf{x}|\mathbf{z})] - \text{KL}(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})), \quad (2)$$

where $p(\mathbf{x}|\mathbf{z})$ is called a *decoder* and $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbf{I})$ is the *prior*. There are various ways of optimizing this lower bound but for continuous \mathbf{z} this could be done efficiently through a re-parameterization of $q(\mathbf{z}|\mathbf{x})$ [9, 15]. Then the architecture is called a *variational auto-encoder* (VAE).

In practice, the inference model assumes a diagonal covariance matrix, *i.e.*, $q(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}(\mathbf{x}), \text{diag}(\boldsymbol{\sigma}^2(\mathbf{x})))$, where $\boldsymbol{\mu}(\mathbf{x})$ and $\boldsymbol{\sigma}^2(\mathbf{x})$ are parameterized by the NN. However, the assumption about the diagonal posterior can be insufficient and not flexible enough to match the true posterior.

2 Improving posterior flexibility using *Normalizing Flows*

A (finite) *normalizing flow*, first formulated by [20, 21] and further developed by [14], is a powerful framework for building flexible posterior distribution by starting with an initial random variable

with a simple distribution for generating $\mathbf{z}^{(0)}$ and then applying a series of invertible transformations $\mathbf{f}^{(t)}$, for $t = 1, \dots, T$. As a result, the last iterate gives a random variable $\mathbf{z}^{(T)}$ that has a more flexible distribution. Once we choose transformations $\mathbf{f}^{(t)}$ for which the Jacobian-determinant can be computed, we aim at optimizing the following objective:

$$\ln p(\mathbf{x}) \geq \mathbb{E}_{q(\mathbf{z}^{(0)}|\mathbf{x})} \left[\ln p(\mathbf{x}|\mathbf{z}^{(T)}) + \sum_{t=1}^T \ln \left| \det \frac{\partial \mathbf{f}^{(t)}}{\partial \mathbf{z}^{(t-1)}} \right| \right] - \text{KL}(q(\mathbf{z}^{(0)}|\mathbf{x})||p(\mathbf{z}^{(T)})). \quad (3)$$

In fact, the normalizing flow can be used to enrich the posterior of the VAE with small or even none modifications in the architecture of the encoder and the decoder.

There are two main kinds of normalizing flows, namely, *general normalizing flows* and *volume preserving flows*. The difference between these types of flow is in the manner how the Jacobian-determinant is handled. The general normalizing flows aim at formulating the flow for which the Jacobian-determinant is relatively easy to compute. On the contrary, the volume-preserving flows design series of transformations such that the Jacobian-determinant equals 1 while still it allows to obtain flexible posterior distributions. The reduced computational complexity is the main reason why volume-preserving flows are so appealing. The question is whether one can propose a series of transformations for which the Jacobian-determinant is equal one, they are cheap to calculate and are general enough to model flexible posteriors. In the next subsection we present a new volume-preserving flow that applies series of Householder transformations that we refer to as the *Householder flow*.

3 Householder Flow

3.1 Motivation

In general, any full-covariance matrix Σ can be represented by the eigenvalue decomposition using eigenvectors and eigenvalues:

$$\Sigma = \mathbf{U} \mathbf{D} \mathbf{U}^\top, \quad (4)$$

where \mathbf{U} is an orthogonal matrix with eigenvectors in columns, \mathbf{D} is a diagonal matrix with eigenvalues. In the case of the vanilla VAE, it would be tempting to model the matrix \mathbf{U} to obtain a full-covariance matrix. The procedure would require a linear transformation of a random variable using an orthogonal matrix \mathbf{U} . Since the absolute value of the Jacobian determinant of an orthogonal matrix is 1, for $\mathbf{z}^{(1)} = \mathbf{U} \mathbf{z}^{(0)}$ one gets $\mathbf{z}^{(1)} \sim \mathcal{N}(\mathbf{U} \boldsymbol{\mu}, \mathbf{U} \text{diag}(\boldsymbol{\sigma}^2) \mathbf{U}^\top)$. If $\text{diag}(\boldsymbol{\sigma}^2)$ coincides with true \mathbf{D} , then it would be possible to resemble the true full-covariance function. Hence, the main goal would be to model the orthogonal matrix of eigenvectors.

Generally, the task of modelling an orthogonal matrix in a principled manner is rather non-trivial. However, first we notice that any orthogonal matrix can be represented in the following form [2, 19]:

Theorem 1. (The Basis-Kernel Representation of Orthogonal Matrices)

For any $M \times M$ orthogonal matrix \mathbf{U} there exist a full-rank $M \times K$ matrix \mathbf{Y} (the basis) and a nonsingular (triangular) $K \times K$ matrix \mathbf{S} (the kernel), $K \leq M$, such that:

$$\mathbf{U} = \mathbf{I} - \mathbf{Y} \mathbf{S} \mathbf{Y}^\top. \quad (5)$$

The value K is called the *degree* of the orthogonal matrix. Further, it can be shown that any orthogonal matrix of degree K can be expressed using the product of Householder transformations [2, 19], namely:

Theorem 2. Any orthogonal matrix with the basis acting on the K -dimensional subspace can be expressed as a product of exactly K Householder transformations:

$$\mathbf{U} = \mathbf{H}_K \mathbf{H}_{K-1} \cdots \mathbf{H}_1, \quad (6)$$

where $\mathbf{H}_k = \mathbf{I} - \mathbf{S}_{kk} \mathbf{Y}_{\cdot k} (\mathbf{Y}_{\cdot k})^\top$, for $k = 1, \dots, K$.

Theoretically, Theorem 2 shows that we can model any orthogonal matrix in a principled fashion using K Householder transformations. Moreover, the Householder matrix \mathbf{H}_k is *orthogonal* matrix itself [7]. Therefore, this property and the Theorem 2 put the Householder transformation as a perfect candidate for formulating a volume-preserving flow that allows to approximate (or even capture) the true full-covariance matrix.

3.2 Definition

The *Householder transformation* is defined as follows. For a given vector $\mathbf{z}^{(t-1)}$ the reflection hyperplane can be defined by a vector (a *Householder vector*) $\mathbf{v}_t(\mathbf{x}) \in \mathbb{R}^M$ that is orthogonal to the hyperplane, and the reflection of this point about the hyperplane is [7]:

$$\mathbf{z}^{(t)} = \left(\mathbf{I} - 2 \frac{\mathbf{v}_t(\mathbf{x})\mathbf{v}_t(\mathbf{x})^\top}{\|\mathbf{v}_t(\mathbf{x})\|^2} \right) \mathbf{z}^{(t-1)} \quad (7)$$

$$= \mathbf{H}_t(\mathbf{x})\mathbf{z}^{(t-1)}, \quad (8)$$

where $\mathbf{H}_t(\mathbf{x}) = \mathbf{I} - 2 \frac{\mathbf{v}_t(\mathbf{x})\mathbf{v}_t(\mathbf{x})^\top}{\|\mathbf{v}_t(\mathbf{x})\|^2}$ is called the *Householder matrix*.

The most important property of $\mathbf{H}_t(\mathbf{x})$ is that it is an orthogonal matrix and hence the absolute value of the Jacobian determinant is equal 1. This fact significantly simplifies the objective (3) because $\ln \left| \det \frac{\partial \mathbf{H}_t(\mathbf{x})\mathbf{z}^{(t-1)}}{\partial \mathbf{z}^{(t-1)}} \right| = 0$, for $t = 1, \dots, T$. Starting from a simple posterior with the diagonal covariance matrix for $\mathbf{z}^{(0)}$, the series of T linear transformations given by (7) defines a new type of volume-preserving flow that we refer to as the *Householder flow* (HF). The vectors $\mathbf{v}_t(\mathbf{x})$, $t = 1, \dots, T$, are produced by the encoder network along with means and variances. The idea of the Householder flow is schematically presented in Figure 1. Once the encoder returns the set of Householder vectors, the Householder flow requires T linear operations to produce a sample from a more flexible posterior with (approximate) full-covariance matrix.




Figure 1: A schematical representation of the encoder network with the Householder flow. (a) The general architecture of the VAE+HF: The encoder returns means and variances for the posterior and a set of Householder vectors that are used to formulate the Householder flow. (b) A single step of the Householder flow that uses linear Householder transformation. In both figures solid lines correspond to the encoder network and the dashed lines are additional quantities required by the HF.

4 Related Work

In [14] invertible linear normalizing flows with known Jacobian determinant were proposed. These are easy to calculate, *i.e.*, the determinant of the Jacobian can be analytically computed, however, many such transformations are needed to capture high-dimensional dependencies. A different approach relies on *volume-preserving flows*, for which the absolute value of Jacobian determinant is equal 1, such as, Non-linear Independent Components Estimation (NICE) [5], Hamiltonian Variational Inference (HVI) [17] and linear Inverse Autoregressive Flow (linIAF) or its non-linear version (nIAF) [10].

The HF is similar in spirit to the linIAF where the linear transformation is also applied but it is the lower triangular inverse Cholesky matrix with ones on the diagonal instead of the Householder matrix. Nevertheless, the motivation of the linIAF differs from ours completely.

The Householder transformations (reflections) were also exploited in the context of learning recurrent neural nets [1]. They were used for modelling unitary weights instead of more flexible variational posterior, however, these served as a component for representing a matrix of eigenvectors, similarly to our approach.

5 Experiments

5.1 MNIST

Set-up In the first experiment we used the MNIST dataset [11] that contains 60,000 training and 10,000 test images of ten handwritten digits (28×28 pixels in size). From the training set we put aside 10,000 images for validation and tuning hyper-parameters. We used the binarized dataset as in [16].

In order to make more reliable comparison to [14, 17], we trained the VAE with 40 stochastic hidden units and the encoder and decoder were parameterized with two-layered neural networks (300 hidden units per layer) and the \tanh activation function. The length of the HF was $T = \{1, 2\}$ (VAE+HF ($T=1/2$)). For training we utilized ADAM [8] with the mini-batch size equal 100 and one example for estimating the expected value. The learning rate was set according to the validation set. The maximum number of epochs was 1500 and early-stopping with a look-ahead of 100 epochs was applied. We used either the *warm-up* [3, 18] for first 200 epochs or *free-bits* constraint [10] with $\lambda \in \{0.25, 0.5\}$, which was determined according to the validation log-probabilities. We initialized weights according to [6].

We compared our approach to linear normalizing flow (VAE+NF) [14], linIAF but with the VAE parameterized by the convolutional network with ResNet blocks (convVAE+linIAF) and its non-linear version (convVAE+nlinIAF) [10], and finite volume-preserving flows: NICE (VAE+NICE) [5], HVI (VAE+HVI) [17]. When appropriate, the length of a flow T is given. Additionally, we provide results for the HVI with the VAE parameterized using convolutional networks (convVAE+HVI) [17] and the VAE (one and two stochastic hidden layers) with Variational Gaussian Process (VAE(1/2hl)+VGP). The methods were compared according to the lower bound of marginal log-likelihood measured on the test set.

Table 1: Comparison of the lower bound of marginal log-likelihood measured in nats of the digits in the MNIST test set. For first three methods the experiment was repeated 3 times.

Method	$\leq \ln p(\mathbf{x})$
VAE	-93.9 ± 0.1
VAE+HF($T=1$)	-80.5 ± 0.3
VAE+HF($T=2$)	-79.5 ± 0.2
VAE+NF ($T=10$) [14]	-87.5
VAE+NF ($T=80$) [14]	-85.1
VAE+NICE ($T=10$) [5]	-88.6
VAE+NICE ($T=80$) [5]	-87.2
VAE+HVI ($T=1$) [17]	-91.7
VAE+HVI ($T=8$) [17]	-88.3
convVAE+HVI ($T=16$) [17]	-84.1
convVAE+linIAF [10]	-84.1 ± 0.1
convVAE+nlinIAF [10]	-80.8 ± 0.1
VAE(1hl)+VGP [22]	-84.8
VAE(2hl)+VGP [22]	-81.3
DRAW+VGP [22]	-79.9




Figure 2: Detailed results for VAE, VAE+HF ($T=1$) and VAE+HF ($T=2$). KL denotes the part of ELBO associated with the Kullback-Leibler Divergence and RE corresponds to the reconstruction error (the expected value of the decoder). The numbers represent negative values of KL and RE.


Discussion The results presented in Table 1 reveal that the proposed flow outperforms other normalizing flows. This outcome is especially interesting since our approach can be easily parallelized and sampling from the variational posterior requires only application of linear transformations, thus, the HF is reasonably cheap computationally. In our current implementation¹ run on GeForce GTX 960M, the wall-clock time for the VAE is about 0.33ms per mini-batch while for the VAE+HF it is about 0.38ms and 0.42ms per mini-batch for T equal 1 and 2, respectively. The average number of epochs until convergence (*i.e.*, excluding the early-stopping) for VAE was 1006 and for VAE+HF with T equal 1 and 2 it was 1042 and 699, respectively. Moreover, in our approach we need to

¹<https://github.com/jmtomczak>


store only $T \times M$ more parameters comparing to the vanilla VAE while, *e.g.*, the VAE+linIAF and VAE+NF requires $\frac{M \times (M-1)}{2}$ and $\mathcal{O}(T \times M)$ more parameters, respectively. All these remarks are in favor of the proposed flow.

In order to get some insight into the behavior of the HF we also examined the values of components of the training objective (see Figure 2). First, the application of the HF helps the VAE to better model data (lower reconstruction error). Second, the HF gives higher flexibility of the posterior because the approximation of the orthogonal matrix allows the initial posterior to model eigenvalues instead of the whole information about the true posterior, which results in smaller Kullback-Leibler penalty (*i.e.*, the variances do not need to model whole information about the true posterior distribution).


We also examined the two-dimensional latent space trained using the vanilla VAE and the VAE+HF (see Figure 3). In case of the vanilla VAE there is a part of the latent space where the generations do not represent digits (see upper right corner in Figure 3(a)). This issue is alleviated by the VAE+HF ($T=1$), however, the most interesting latent representation is obtained by applying two Householder transformations (see Figure 3(b) and (c), respectively).



(a) VAE



(b) VAE+HF ($T=1$)



(c) VAE+HF ($T=2$)

Figure 3: Visualization of the two-dimensional latent space of a generative model for (a) the vanilla VAE, and the VAE+HF with (b) $T=1$ and (c) $T=2$.

5.2 Histopathology data

Data preparation In the second experiment we applied the VAE and the VAE+HF to a more challenging problem of grayscale histopathology data. We prepared a dataset basing on histopathological images freely available on-line². We selected 16 patients³ and each histopathological image represented a bone marrow biopsy. Diagnoses of the chosen cases were associated with different kinds of cancer (*e.g.*, lymphoma, leukemia) or anemia. All images were taken using HE, 40 \times , and

²<http://www.enjopath.com/>
³Patient numbers: 272, 274, 283, 289, 290, 291, 292, 295, 297, 298, 299.

each image was of size 336×448 . The original RGB representation was transformed to grayscale⁴. Further, we divided each image into small patches of size 28×28 (see Figure 4 for exemplary image patches). Eventually, we picked 10 patients for training, 3 patients for validation and 3 patients for testing, which resulted in 6,800 training images, 2,000 validation images and 2,000 test images. The selection of patients was performed in such fashion that each dataset contained representative images with different diagnoses and amount of fat.




Figure 4: Selected image patches from the prepared histopathological data.

Set-up We used the same architecture of the VAE and the same training procedure as in the previous experiment. Since the data are grayscale, we modeled data using normal distribution instead of the Bernoulli distribution with means constrained to $[0, 1]$ by applying the sigmoid function. The only difference in the training procedure was that we performed model selection for *free-bits* with $\lambda \in \{0.25, 0.5, 0.75, 1.0, 1.5, 2.0\}$. We compared the vanilla VAE with the VAE+HF ($T=1,2$).

Method	$\leq \ln p(\mathbf{x})$
VAE	1033.33 ± 2.14
VAE+HF ($T=1$)	1041.13 ± 1.72
VAE+HF ($T=2$)	1032.29 ± 2.13

Table 2: Comparison of the lower bound of marginal log-likelihood measured in nats of the image patches in the Histopathology test set. The experiment was repeated 5 times.

Discussion The results presented in Table 2 confirmed findings of the previous experiment that the application of the HF helps in obtaining more flexible posterior, however, a series of two Householder transformations performed similarly to the vanilla VAE. A possible explanation for this effect is that the application of two Householder transformations resulted in larger and more expressive model that was harder to learn using a relatively small sample size.

6 Conclusion

In this paper, we proposed a new volume-preserving flow using Householder transformations. Our idea relies on the observation that the true full-covariance matrix can be decomposed to the diagonal matrix with eigenvalues on the diagonal and the orthogonal matrix of eigenvalues that could be further modeled using a series of Householder transformations. The obtained results in the experimental studies reveal that fully connected VAEs with Householder flows could perform better than their convolutional counterparts. This makes the HF a very promising direction for further investigation. In the future, we aim at using the outlined flow in recently proposed extensions of the variational inference, *e.g.*, importance weighted VAE [4], Rényi Divergence for VAE [12] and Ladder VAE [18]. Moreover, we believe that the proposed flow can be beneficiary in modelling natural and medical images, where flexibility of the posterior is crucial [13, 18]. We leave investigating this issue for future research.

⁴A pixel’s value was determined according to the formula: $0.299 \times R + 0.587 \times G + 0.114 \times B$.

Acknowledgments

The research conducted by Jakub M. Tomczak was funded by the European Commission within the Marie Skłodowska-Curie Individual Fellowship (Grant No. 702666, "Deep learning and Bayesian inference for medical imaging").

References

- [1] Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. *ICML*, 2016.
- [2] Christian Bischof and Xiaobai Sun. On orthogonal block elimination. *Argonne National Laboratory, Argonne, IL, Tech. Rep. MCS-P450-0794*, 1994.
- [3] 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.
- [4] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- [5] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [6] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, volume 9, pages 249–256, 2010.
- [7] Alston S Householder. Unitary triangularization of a nonsymmetric matrix. *Journal of the ACM (JACM)*, 5(4):339–342, 1958.
- [8] Diederik Kingma and Jimmy Ba. ADAM: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [9] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [10] Diederik P Kingma, Tim Salimans, Rafal Józefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improving variational inference with inverse autoregressive flow. *NIPS*, 2016.
- [11] Yann LeCun, Corinna Cortes, and Christopher JC Burges. The MNIST database of handwritten digits, 1998.
- [12] Yingzhen Li and Richard E Turner. Rényi divergence variational inference. *arXiv preprint arXiv:1602.02311*, 2016.
- [13] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *ICML*, pages 1747–1756, 2016.
- [14] Danilo Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. In *ICML*, pages 1530–1538, 2015.
- [15] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [16] Ruslan Salakhutdinov and Iain Murray. On the quantitative analysis of deep belief networks. In *ICML*, pages 872–879. ACM, 2008.
- [17] Tim Salimans, Diederik P Kingma, and Max Welling. Markov chain Monte Carlo and Variational Inference: Bridging the gap. In *ICML*, pages 1218–1226, 2015.
- [18] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder variational autoencoders. *arXiv preprint arXiv:1602.02282*, 2016.
- [19] Xiaobai Sun and Christian Bischof. A basis-kernel representation of orthogonal matrices. *SIAM Journal on Matrix Analysis and Applications*, 16(4):1184–1196, 1995.
- [20] EG Tabak and Cristina V Turner. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- [21] Esteban G Tabak, Eric Vanden-Eijnden, et al. Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233, 2010.
- [22] Dustin Tran, Rajesh Ranganath, and David M Blei. The Variational Gaussian Process. *ICLR*, 1050:23, 2016.