Learning gradient-based ICA
by neurally estimating mutual information

Hlynur Davíð Hlynsson, Laurenz Wiskott

February 2019

Abstract

Several methods of estimating the mutual information of random variables have been developed in recent years. They can prove valuable for novel approaches to learning statistically independent features. In this paper, we use one of these methods, a mutual information neural estimation (MINE) network, to present a proof-of-concept of how a neural network can perform linear ICA. We minimize the mutual information, as estimated by a MINE network, between the output units of a differentiable encoder network. This is done by simple alternate optimization of the two networks. The method is shown to get a qualitatively equal solution to FastICA on blind-source-separation of noisy sources.

1 Introduction

Independent component analysis (ICA) aims at estimating unknown *sources* that have been mixed together into an *observation*. The usual assumptions are that the sources are statistically independent and no more than one is Gaussian [1]. The now-cemented metaphor is one of a cocktail party problem: several people (sources) are speaking simultaneously and their speech has been mixed together in a recording (observation). The task is to unmix the recording such that all dialogues can be listened to clearly.

In linear ICA, we have a data matrix S whose rows are drawn from statistically independent distributions, a mixing matrix A, and an observation matrix X:

$$X = AS$$

and we want to find an unmixing matrix U of A that recovers the sources up to a permutation and scaling:

$$Y = UX$$

The general non-linear ICA problem is ill-posed [2, 3] as there is an infinite number of solutions if the space of mixing functions is unconstrained. However, post-linear [4] (PNL) ICA is solvable. This is a particular case of non-linear ICA where the observations take the form

$$X = f(AS)$$

where f operates componentwise, i.e. $X_{i,t} = f_i \left(\sum_n^N A_{i,n} S_{n,t} \right)$. The problem is solved efficiently if f is at least approximately invertible [5] and there are approaches to optimize the problem for non-invertible f as well [6]. For signals with time-structure, however, the problem is not ill-posed even though it is for i.i.d. samples [7, 8].

To frame ICA as an optimization problem, we must find a way to measure the statistical independence of the output components and minimize this quantity. There are two main ways to approach this: either minimize the mutual information between the sources [9, 10, 11], or maximize the sources' non-Gaussianity [12, 13].

There has been a recent interest in combining deep learning with the principles of ICA, usually in an adversarial framework, for example Deep InfoMax (DIM) [14], Graph Deep InfoMax [15] and Generative adversarial networks [16], which utilize the work of Brakel et al. [17]. Our work is distinct from theirs as we do not rely on adversarial training.

2 Method

We train an encoder E to generate an output $(z_1, z_2, \ldots, z_M)$ such that any one of the output components is statistically independent of the union of the others, i.e. $P(z_i, \boldsymbol{z_{-i}}) = P(z_i)P(\boldsymbol{z_{-i}})$, where

$$\boldsymbol{z_{-i}} := (z_1, \ldots, z_{i-1}, z_{i+1}, \ldots, z_M)$$

The statistical independence of z_i and $\boldsymbol{z_{-i}}$ can be maximized by minimizing their mutual information

$$I\left(Z_i; \boldsymbol{Z_{-i}}\right) = \int_z \int_{\boldsymbol{z_{-i}}} p(z_i, \boldsymbol{z_{-i}}) \log \left(\frac{p(z_i, \boldsymbol{z_{-i}})}{p(z_i)p(\boldsymbol{z_{-i}})} \right) dz_i d\boldsymbol{z_{-i}} \qquad (1)$$

This quantity is hard to estimate, particularly for high-dimensional data. We therefore estimate the lower bound of Eq. (1) using a mutual information neural estimation (MINE) network M [18]:

$$I\left(Z_i; \boldsymbol{Z_{-i}}\right) \geq L_i = \mathbb{E}_{\mathbb{J}}\left[M\left(z_i, \boldsymbol{z_{-i}}\right)\right] - \log \left(\mathbb{E}_{\mathbb{M}}\left[e^{M\left(z_i, \boldsymbol{z_{-i}}\right)} \right] \right) \qquad (2)$$

where $\mathbb{J}$ indicates that the expected value is taken over the joint and similarly $\mathbb{M}$ for the product of marginals. The networks E and M are parameterized by θ_E and θ_M. The encoder takes the observations as input and the MINE network takes the output of the encoder as an input.

The E network minimizes $L := \sum_i L_i$ in order for the outputs to have low mutual information and therefore be statistically independent. In order to get a faithful estimation of the lower bound of the mutual information, the M network maximizes L. Thus, in a push-pull fashion the system as a whole converges to independent output components of the encoder network E. In practice, rather than training the E and M networks simultaneously it proved useful to train M from scratch for a

few iterations after each iteration of training E, since the loss functions of E and M are at odds with each other. When the encoder is trained, the MINE network's parameters are frozen and *vice versa.*
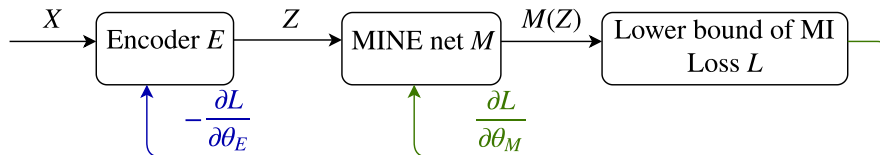

Figure 1: The system learns statistically independent outputs by alternate optimization of an encoder E and a MINE network M parameterized by θ_E and θ_M. The MINE objective (Eq. 2) is minimized with respect to θ_E for weight updates of the encoder but it is *maximized* with respect to θ_M for weight updates of the MINE network.

3   Results

We validate the method[1] for linear noisy ICA example [19]. Three independent, noisy sources — sine wave, square wave and sawtooth signal (Fig. 2a) — are mixed linearly (Fig. 2b):

$$Y = \begin{bmatrix} 1 & 1 & 1 \\ 0.5 & 2 & 1 \\ 1.5 & 1 & 2 \end{bmatrix} S$$

The encoder is a single layer neural network with linear activation with a differentiable whitening layer [20] before the output. The whitening layer is a key component for performing successful blind source separation for our method. Statistically independent random variables are necessarily uncorrelated, so whitening the output by construction beforehand simplifies the optimization problem significantly.

The MINE network M is a seven-layer neural network. Each layer but the last one has 64 units with a rectified linear activation function. Each training epoch of the encoder is followed by seven training epochs of M. Estimating the exact mutual information is not essential, so few iterations suffice for a good gradient direction.

Since the MINE network is applied to each component individually, to estimate mutual information (Eq. 2), we need to pass each sample through the MINE network N times — once for each component. Equivalently, one could conceptualize this as having N copies of the MINE network and feeding the samples to it in parallel, with different components singled out. Thus, for sample $(z_1, z_2, \ldots, z_N)$ we feed in $(z_i; z_{-i})$, for each i. Both

---

[1]Full code for the results is available at `github.com/wiskott-lab/gradient-based-ica/blob/master/bss3.ipynb`

(a) The original sources.

(b) Linear mixture of sources.

(c) Sources recovered by our method.
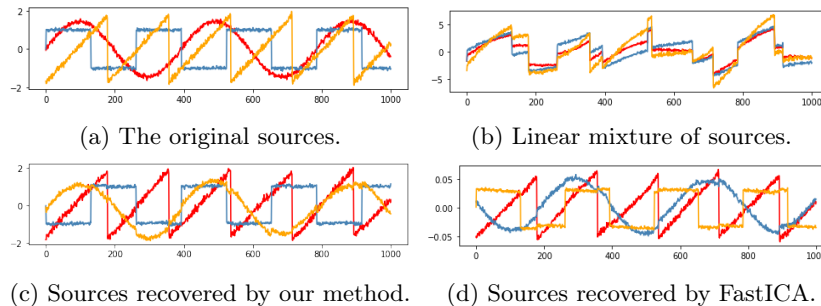
(d) Sources recovered by FastICA.

Figure 2: Three independent, noisy sources (a) are mixed linearly (b). Our method recovers them (c) to the same extent as FastICA (d).

networks are optimized using Nesterov momentum ADAM [21] with a learning rate of 0.005.

For this simple example, our method (Fig. 2c) is equivalently good at unmixing the signals as FastICA (Fig. 2d), albeit slower. Note that, in general, the sources can only be recovered up to permutation and scaling.

4   Summary

We've introduced a proof-of-concept implementation for training a differentiable function for performing ICA. The method consists of alternating the optimization of an encoder and a neural mutual-information neural estimation (MINE) network. The mutual information estimate between each encoder output and the union of the others is minimized with respect to the encoder's parameters. Although this work is in a very preliminary stage, further investigation into the method is warranted. The general nonlinear ICA problem is ill-posed, but it is an interesting question whether this method can work for non-linear problems with low complexity. We can constrain the expresiveness of our encoder by limiting for example the number of layers or number of hidden units in the neural network, thus constraining the solution space of the method. The method is also trivially extended for over- or undercomplete ICA by changing the number of output units. Higher dimensional and real-world data can also be tested.

As this method can be used for general neural network training, it should be investigated whether useful representations can be learned while solving the ICA task. This method blends nicely into deep learning architectures and the MINE loss term can be added as a regularizer to other loss functions. We imagine that this can be helpful for methods such as deep sparse coding to enforce independence between features and disentangle factors of variation.

4

References

[1] C. Jutten and J. Karhunen, "Advances in nonlinear blind source separation," in *Proc. of the 4th Int. Symp. on Independent Component Analysis and Blind Signal Separation (ICA2003)*, pp. 245–256, 2003.

[2] A. Hyvärinen and P. Pajunen, "Nonlinear independent component analysis: Existence and uniqueness results," *Neural Networks*, vol. 12, no. 3, pp. 429–439, 1999.

[3] G. Darmois, "Analyse générale des liaisons stochastiques: etude particulière de l'analyse factorielle linéaire," *Revue de l'Institut international de statistique*, pp. 2–8, 1953.

[4] A. Taleb and C. Jutten, "Source separation in post-nonlinear mixtures," *IEEE Transactions on signal Processing*, vol. 47, no. 10, pp. 2807–2820, 1999.

[5] A. Ziehe, M. Kawanabe, S. Harmeling, and K.-R. Müller, "Blind separation of post-nonlinear mixtures using linearizing transformations and temporal decorrelation," *Journal of Machine Learning Research*, vol. 4, no. Dec, pp. 1319–1338, 2003.

[6] A. Ilin and A. Honkela, "Post-nonlinear independent component analysis by variational Bayesian learning," in *International Conference on Independent Component Analysis and Signal Separation*, pp. 766–773, Springer, 2004.

[7] T. Blaschke, T. Zito, and L. Wiskott, "Independent slow feature analysis and nonlinear blind source separation," *Neural computation*, vol. 19, no. 4, pp. 994–1021, 2007.

[8] H. Sprekeler, T. Zito, and L. Wiskott, "An extension of slow feature analysis for nonlinear blind source separation," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 921–947, 2014.

[9] S.-i. Amari, A. Cichocki, and H. H. Yang, "A new learning algorithm for blind signal separation," in *Advances in neural information processing systems*, pp. 757–763, 1996.

[10] A. J. Bell and T. J. Sejnowski, "A non-linear information maximisation algorithm that performs blind separation," in *Advances in neural information processing systems*, pp. 467–474, 1995.

[11] J.-F. Cardoso, "InfoMax and maximum likelihood for blind source separation," *IEEE Signal processing letters*, vol. 4, no. 4, pp. 112–114, 1997.

[12] A. Hyvärinen and E. Oja, "Independent component analysis: algorithms and applications," *Neural networks*, vol. 13, no. 4-5, pp. 411–430, 2000.

[13] T. Blaschke and L. Wiskott, "CuBICA: Independent component analysis by simultaneous third-and fourth-order cumulant diagonalization," *IEEE Transactions on Signal Processing*, vol. 52, no. 5, pp. 1250–1256, 2004.

[14] R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, A. Trischler, and Y. Bengio, "Learning deep representations by

mutual information estimation and maximization," *arXiv preprint arXiv:R1808.06670*, 2018.

[15] P. Veličković, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph InfoMax," *arXiv preprint arXiv:1809.10341*, 2018.

[16] 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*, pp. 2672–2680, 2014.

[17] P. Brakel and Y. Bengio, "Learning independent features with adversarial nets for non-linear ICA," *arXiv preprint arXiv:1710.05050*, 2017.

[18] M. I. Belghazi, A. Baratin, S. Rajeswar, S. Ozair, Y. Bengio, A. Courville, and R. D. Hjelm, "MINE: Mutual Information Neural Estimation," *arXiv preprint arXiv:1801.04062*, 2018.

[19] "Blind source separation using FastICA." `https://scikit-learn.org/stable/auto_examples/decomposition/plot_ica_blind_source_separation.html`. Accessed: 2019-02-24.

[20] M. Schüler, H. D. Hlynsson, and L. Wiskott, "Gradient-based training of slow feature analysis by differentiable approximate whitening," *arXiv preprint arXiv:1808.08833*, 2018.

[21] T. Dozat, "Incorporating Nesterov momentum into ADAM," 2016.