

# Towards cross-lingual distributed representations without parallel text trained with adversarial autoencoders

**Antonio Valerio Miceli Barone**  
The University of Edinburgh  
Informatics Forum, 10 Crichton Street  
Edinburgh  
amiceli@inf.ed.ac.uk

## Abstract

Current approaches to learning vector representations of text that are compatible between different languages usually require some amount of parallel text, aligned at word, sentence or at least document level. We hypothesize however, that different natural languages share enough semantic structure that it should be possible, in principle, to learn compatible vector representations just by analyzing the monolingual distribution of words.

In order to evaluate this hypothesis, we propose a scheme to map word vectors trained on a source language to vectors semantically compatible with word vectors trained on a target language using an adversarial autoencoder.

We present preliminary qualitative results and discuss possible future developments of this technique, such as applications to cross-lingual sentence representations.

## 1 Introduction

Distributed representations that map words, sentences, paragraphs or documents to vectors real numbers have proven extremely useful for a variety of natural language processing tasks (Bengio et al., 2006; Collobert and Weston, 2008; Turian et al., 2010; Maas et al., 2011; Mikolov et al., 2013b; Socher et al., 2013; Pennington et al., 2014; Levy and Goldberg, 2014; Le and Mikolov, 2014; Baroni et al., 2014; Levy et al., 2015), as they provide an effective way to inject into machine learning models general prior knowledge about language automatically obtained from inexpensive unannotated corpora. Based on the assumption that different languages share a similar semantic struc-

ture, various approaches succeeded to obtain distributed representations that are compatible across multiple languages, either by learning mappings between different embedding spaces (Mikolov et al., 2013a; Faruqui and Dyer, 2014) or by jointly training cross-lingual representations (Klementiev et al., 2012; Hermann and Blunsom, 2013; Chandar et al., 2014; Gouws et al., 2014). These approaches all require some amount of parallel text, aligned at word level, sentence level or at least document level, or some other kind of parallel resources such as dictionaries (Ammar et al., 2016).

In this work we explore whether the assumption of a shared semantic structure between languages is strong enough that it allows to induce compatible distributed representations without using any parallel resource. We only require monolingual corpora that are thematically similar between languages in a general sense.

We hypothesize there exist a suitable vectorial space such that each language can be viewed as a random process that produces vectors at some level of granularity (words, sentences, paragraphs, documents) which are then encoded as discrete surface forms, and we hypothesize that, if languages are used to convey thematically similar information in similar contexts, these random processes should be approximately isomorphic between languages, and that this isomorphism can be learned from the statistics of the realizations of these processes, the monolingual corpora, in principle without any form of explicit alignment.

We motivate this hypothesis by observing that humans, especially young children, who acquire multiple languages, can often do so with relatively little exposure to explicitly aligned parallel linguistic information, at best they may have access to distant and noisy alignment information in the form of multisensorial environmental clues. Nevertheless, multilingual speakers are always au-

tomatically able to translate between all the languages that they can speak, which suggests that their brain either uses a shared conceptual representations for the different surface features of each language, or uses distinct but near-isomorphic representations that can be easily transformed into each other.

## 2 Learning word embedding cross-lingual mappings with adversarial autoencoders

The problem of learning transformations between probability distributions of real vectors has been studied in the context of generative neural network models, with approaches such as Generative Moment Matching Networks (GMMNs) (Li et al., 2015) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014). In this work we consider GANs, since their effectiveness has been demonstrated in the literature more thoroughly than GMMNs.

In a typical GAN, we wish to train a *generator* model, usually a neural network, to transform samples from a known, easy to sample, uninformative distribution (e.g. Gaussian or uniform) into samples distributed according to a target distribution defined implicitly by a training set. In order to do so, we iteratively alternate between training a differentiable *discriminator* model, also a neural network, to distinguish between training samples and artificial samples produced by the generator, and training the generator to fool the discriminator into misclassifying the artificial examples as training examples. This can be done with conventional gradient-based optimization because the discriminator is differentiable thus it can backpropagate gradients into the generator.

It can be proven that, with sufficient model capacity and optimization power, sufficient entropy (information dimension) of the generator input distribution, and in the limit of infinite training set size, the generator learns to produce samples from the correct distribution. Intuitively, if there is any computable test that allows to distinguish the artificial samples from the training samples with better than random guessing probability, then a sufficiently powerful discriminator will eventually learn to exploit it and then a sufficiently powerful generator will eventually learn to counter it, until the generator output distribution becomes undistinguishable from the true training

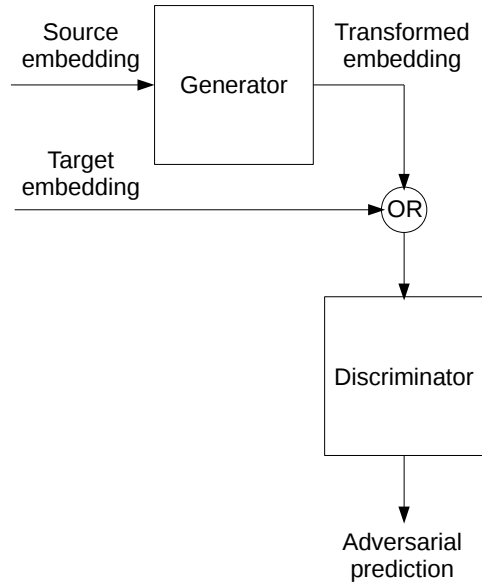


Figure 1: Generative adversarial network for cross-lingual embedding mapping

distribution. In practice, actual models have finite capacity and gradient-based optimization algorithms can become unstable or stuck when applied to this multi-objective optimization problem, though they have been successfully used to generate fairly realistic-looking images (Denton et al., 2015; Radford et al., 2015).

In our preliminary experiments we attempted to adapt GANs to our problem, by training the generator to learn a transformation between word embeddings trained on different languages (fig. 1). Let  $d$  be the embedding dimensionality,  $G_{\theta_G} : \mathcal{R}^d \rightarrow \mathcal{R}^d$  be the generator parametrized by  $\theta_G$ ,  $D_{\theta_D} : \mathcal{R}^d \rightarrow [0, 1]$  be the discriminator parametrized by  $\theta_D$ .

At each training step:

1. draw a sample  $\{f\}_n$  of  $n$  source embeddings, according to their (adjusted) word frequencies
2. transform them into target-like embeddings  $\{\hat{e}\}_n = G_{\theta_G}(\{f\}_n)$
3. evaluate them with the discriminator, estimating their probability of having been sampled from the true target distribution  $\{p\}_n = D_{\theta_D}(\{\hat{e}\})$
4. update the generator parameters  $\theta_G$  to reduce the average adversarial loss  $L_a = -\log(\{p\}_n)$

5. draw a sample  $\{e\}_n$  of  $n$  true target embeddings
6. update the discriminator parameters  $\theta_D$  to reduce its binary cross-entropy loss on the classification between  $\{e\}_n$  (positive class) and  $\{\hat{e}\}$  (negative class)

repeat these steps until convergence.

Unfortunately we found that in this setup, even with different network architectures and hyperparameters, the model quickly converges to a pathological solution where the generator always emits constant or near-constant samples that somehow can fool the discriminator. This appears to be an extreme case of the know mode-seeking issue of GANs (Radford et al., 2015; Theis et al., 2015; Salimans et al., 2016), which is probably exacerbated in our settings because of the point-mass nature of our probability distributions where each word embedding is a mode on its own.

In order to avoid these pathological solutions, we needed a way to penalize the generator for destroying too much information about its input. Therefore we turned our attention to Adversarial Autoencoders (AAE) (Makhzani et al., 2015). In an AAE, the generator, now called *encoder*, is paired with another model, the *decoder*  $R_{\theta_R} : \mathcal{R}^d \rightarrow \mathcal{R}^d$  parametrized by  $\theta_R$  which attempts to transform the artificial samples emitted by the encoder back into the input samples. The encoder and the decoder are jointly trained to minimize a combination of the average reconstruction loss  $L_r(\{f\}_n, R_{\theta_R}(G_{\theta_G}(\{f\}_n)))$  and the adversarial loss defined as above. The discriminator is trained as above. In the original formulation of the AAE, the discriminator is used to enforce a known prior (e.g. Gaussian or Gaussian mixture) on the intermediate, *latent* representation, in our setting instead we use it to match the latent representation to the target embedding distribution so that the encoder can be used to transform source embeddings into target ones (fig. 2).

In our experiments, we use the cosine dissimilarity as reconstruction loss, and as a further penalty we also include the pairwise cosine dissimilarity between the generated latent samples  $\{\hat{e}\}$  and the true target samples  $\{e\}_n$ . Therefore, the total loss incurred by the encoder-decoder at each step is

$$L_{GR} = \lambda_r L_r(\{f\}_n, R_{\theta_R}(G_{\theta_G}(\{f\}_n))) - \lambda_a \log(\{p\}_n) + \lambda_c L_r(\{e\}_n, G_{\theta_G}(\{f\}_n))$$

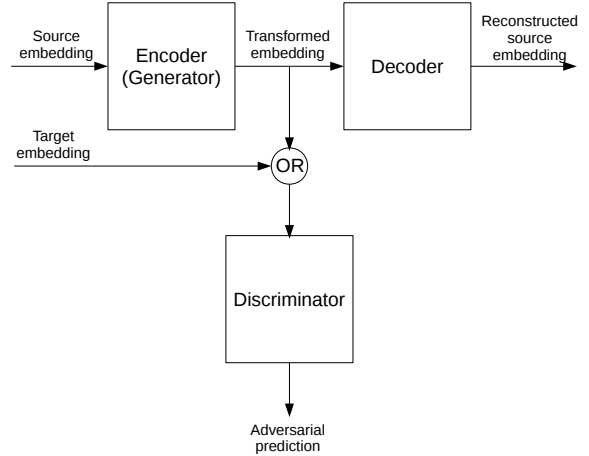


Figure 2: Adversarial autoencoder for cross-lingual embedding mapping (loss function blocks not shown).

where  $\lambda_r$ ,  $\lambda_a$  and  $\lambda_c$  are hyperparameters (all set equal to 1 in our experiments).

### 3 Experiments

We performed some preliminary exploratory experiments on our model. In this section we report salient results.

The first experiment is qualitative, to assess whether our model is able to learn any semantically sensible transformation at all. We consider English to Italian embedding mapping.

We train English and Italian word embeddings on randomly subsampled Wikipedia corpora consisting of about 1.5 million sentences per language. We use word2vec (Mikolov et al., 2013b) in skipgram mode to generate embeddings with dimension  $d = 100$ . Our encoder and decoder are linear models with tied matrices (one the transpose of the other), initialized as random orthogonal matrices (we also explored deep non-linear autoencoders but we found that they make the optimization more difficult without providing apparent benefits).

Our discriminator is a Residual Network (He et al., 2015) without convolutions, one leaky ReLU non-linearity (Maas et al., 2013) per block, no non-linearities on the passthrough path, batch normalization (Ioffe and Szegedy, 2015) and dropout (Nitish et al., 2014). The block (layer) equation is:

$$h_{t+1} = \phi(W_t \times h_{t-1}) + h_{t-1} \quad (1)$$

where  $W_t$  is a weight matrix and  $\phi$  is batch normalization (with its internal parameters) followed

by leaky ReLU and  $h_t$  is a  $k$ -dimensional block state (in our experiments  $k = 40$ ). The network has  $T = 10$  blocks followed by a 1-dimensional output layer with logistic sigmoid activation. We found that using a Residual Network as discriminator rather than a standard multi-layer perceptron yields larger gradients being backpropagated to the generator, facilitating training. We actually train two discriminators per experiment, with identical structure but different random initializations, and use one to train the generator and the other for monitoring in order to help us determine whether overfitting or underfitting occurs.

At each step, word embeddings are sampled according to their frequency in the original corpora, adjusted to subsample frequent words, as in `word2vec`. Updates are performed using the Adam optimizer (Kingma and Ba, 2014) with learning rate 0.001 for the encoder-decoder and 0.01 for the discriminator.

The code<sup>1</sup> is implemented in Python, Theano (Theano Development Team, 2016) and Lasagne.

We qualitatively analyzed the quality of the embeddings by considering the closest Italian embeddings to a sample of transformed English embeddings. We notice that in some cases the closest or nearly closest embedding is the true translation, for instance 'computer' (en)  $\rightarrow$  'computer' (it). In other cases, the closest terms are not translations but subjectively appear to be semantically related, for instance 'rain' (en)  $\rightarrow$  'gelo', 'gela', 'intensissimo', 'galleggiava', 'rigidissimo', 'arida', 'semi-desertico', 'fortunale', 'gelata', 'piovosa' (it 10-best), or 'comics' (en)  $\rightarrow$  'Kadath', 'Microciccio', 'Cugel', 'Promethea', 'flashback', 'episodio', 'Morimura', 'Chatwin', 'romanzato', 'Deedlit' (it 10-best), or 'anime' (en)  $\rightarrow$  'Zatanna', 'Alita', 'Yuriko', 'Wildfire', 'Carmilla', 'Batwoman', 'Leery', 'Aquarion', 'Vampirella', 'Minaccia' (it 10-best). Other terms, such as names of places however, tend to be transformed incorrectly, for instance 'France' (en)  $\rightarrow$  'Radiomobile', 'Cartubi', 'Freniatrria', 'UNUCI', 'Cornhole', 'Internazione', 'CSCE', 'Folklorica', 'UECI', 'Rientro' (it 10-best).

We further evaluate our model on German to English and English to German embedding transformations, using the same evaluation setup as (Klementiev et al., 2012) with embeddings trained

on the concatenation of the Reuters corpora and the News Commentary 2015 corpora, with embedding dimension  $d = 40$  and discriminator depth  $T = 4$ . On a qualitative analysis notice similar partial semantic similarity patterns. However the cross-lingual document classification task we were able to improve over the baseline only for the smallest training set size.

## 4 Discussion and future work

From the qualitative analysis of the word embedding mappings it appears that the model does learn to transfer some semantic information, although it's not competitive with other cross-lingual representation approaches. This may be possibly an issue of hyperparameter choice and architectural details, since, to our knowledge, this is the first work to apply adversarial training techniques to point-mass distribution arising from NLP tasks.

Further experimentation is needed to determine whether the model can be improved or whether we already hit a fundamental limit on how much semantic transfer can be performed by monolingual distribution matching alone. This additional experimentation may help to test how strongly our initial hypothesis of semantic isomorphism between languages holds, in particular across languages of different linguistic families.

Even if this hypothesis does not hold in a strong sense and semantic transfer by monolingual text alone turns out to be infeasible, our technique might help in conjunction with training on parallel data. For instance, in neural machine translation "sequence2sequence" transducers without attention (Cho et al., 2014), it could be useful to train as usual on parallel sentences and train in autoencoder mode on monolingual sentences, using an adversarial loss computed by a discriminator on the intermediate latent representations to push them to be isomorphic between languages. A modification of this technique that allows for the latent representation to be variable-sized could be also applied to the attentive "sequence2sequence" transducers (Bahdanau et al., 2014), as an alternative or in addition to monolingual dataset augmentation by backtranslation (Sennrich et al., 2015).

Furthermore, it may be worth to evaluate additional distribution learning approaches such as the aforementioned GMMs, as well as the more recent BiGAN/ALI framework (Donahue et al., 2016; Dumoulin et al., 2016) which uses an adver-

<sup>1</sup>Code with full hyperparameters available at: <https://github.com/Avmb/clweadv>

sarial discriminator loss both to match latent distributions and to enforce reconstruction, and also to consider more recent GAN training techniques (Salimans et al., 2016).

In conclusion we believe that this work initiates a potentially promising line of research in natural language processing consisting of applying distribution matching techniques such as adversarial training to learn isomorphisms between languages.

## Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement 645452 (QT21).

## References

- Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A Smith. 2016. Massively multilingual word embeddings. *arXiv preprint arXiv:1602.01925*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *ACL (1)*, pages 238–247.
- Yoshua Bengio, Holger Schwenk, Jean-Sébastien Senécal, Frédéric Morin, and Jean-Luc Gauvain. 2006. Neural probabilistic language models. In *Innovations in Machine Learning*, pages 137–186. Springer.
- Sarath Chandar, Stanislas Lauly, Hugo Larochelle, Mitesh Khapra, Balaraman Ravindran, Vikas C Raykar, and Amrita Saha. 2014. An autoencoder approach to learning bilingual word representations. In *Advances in Neural Information Processing Systems*, pages 1853–1861.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167. ACM.
- Emily L Denton, Soumith Chintala, Rob Fergus, et al. 2015. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in Neural Information Processing Systems*, pages 1486–1494.
- J. Donahue, P. Krähenbühl, and T. Darrell. 2016. Adversarial Feature Learning. *ArXiv e-prints*, May.
- V. Dumoulin, I. Belghazi, B. Poole, A. Lamb, M. Arjovsky, O. Mastropietro, and A. Courville. 2016. Adversarially Learned Inference. *ArXiv e-prints*, June.
- Manaal Faruqi and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. Association for Computational Linguistics.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2014. Bilbowa: Fast bilingual distributed representations without word alignments. *arXiv preprint arXiv:1410.2455*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.
- Karl Moritz Hermann and Phil Blunsom. 2013. Multilingual distributed representations without word alignment. *arXiv preprint arXiv:1312.6173*.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing crosslingual distributed representations of words. In *Proceedings of the International Conference on Computational Linguistics (COLING)*, Bombay, India, December.
- Quoc V Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. *arXiv preprint arXiv:1405.4053*.
- Omer Levy and Yoav Goldberg. 2014. Neural word embedding as implicit matrix factorization. In *Advances in Neural Information Processing Systems*, pages 2177–2185.
- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225.

- Yujia Li, Kevin Swersky, and Richard Zemel. 2015. Generative moment matching networks. *arXiv preprint arXiv:1502.02761*.
- Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 142–150. Association for Computational Linguistics.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30, page 1.
- Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, and Ian Goodfellow. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013a. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Srivastava Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543.
- Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. 2016. Improved Techniques for Training GANs. *ArXiv e-prints*, June.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Richard Socher, John Bauer, Christopher D Manning, and Andrew Y Ng. 2013. Parsing with compositional vector grammars. In *ACL (1)*, pages 455–465.
- Theano Development Team. 2016. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May.
- Lucas Theis, Aäron van den Oord, and Matthias Bethge. 2015. A note on the evaluation of generative models. *arXiv preprint arXiv:1511.01844*.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pages 384–394. Association for Computational Linguistics.