

Learning Robust Representations of Text

Yitong Li Trevor Cohn Timothy Baldwin

Department of Computing and Information Systems

The University of Melbourne, Australia

yitongl4@student.unimelb.edu.au, {tcohn,tbaldwin}@unimelb.edu.au

Abstract

Deep neural networks have achieved remarkable results across many language processing tasks, however these methods are highly sensitive to noise and adversarial attacks. We present a regularization based method for limiting network sensitivity to its inputs, inspired by ideas from computer vision, thus learning models that are more robust. Empirical evaluation over a range of sentiment datasets with a convolutional neural network shows that, compared to a baseline model and the dropout method, our method achieves superior performance over noisy inputs and out-of-domain data.¹

1 Introduction

Deep learning has achieved state-of-the-art results across a range of computer vision (Krizhevsky et al., 2012), speech recognition (Graves et al., 2013) and natural language processing tasks (Bahdanau et al., 2015; Kalchbrenner et al., 2014; Yih et al., 2014; Bitvai and Cohn, 2015). However, deep models are often overconfident for noisy test instances, making them susceptible to adversarial attacks (Nguyen et al., 2015; Tabacof and Valle, 2016). Goodfellow et al. (2014) argued that the primary cause of neural networks' vulnerability to adversarial perturbation is their linear nature, due to neural models being intentionally designed to behave in a mostly linear manner to facilitate optimization. Fawzi et al. (2015) provided a theoretical framework for analyzing the

robustness of classifiers to adversarial perturbations, and also showed linear models are usually not robust to adversarial noise.

In this work, we present a regularization method which makes deep learning models more robust to noise, inspired by Rifai et al. (2011). The intuition behind the approach is to stabilize predictions by minimizing the ability of features to perturb predictions, based on high-order derivatives. Rifai et al. (2011) introduced contractive auto-encoders based on similar ideas, using the Frobenius norm of the Jacobian matrix as a penalty term to extract robust features. Further, Gu and Rigazio (2014) introduced deep contractive networks, generalizing this idea to a feed-forward neural network. Also related, Martens (2010) investigated a second-order optimization method based on Hessian-free approach for training deep auto-encoders. Where our proposed approach differs is that we train models using first-order derivatives of the training loss as part of a regularization term, necessitating second-order derivatives for computing the gradient. We empirically demonstrate the effectiveness of the model over text corpora with increasing amounts of artificial masking noise, using a range of sentiment analysis datasets (Pang and Lee, 2008) with a convolutional neural network model (Kim, 2014). In this, we show that our method is superior to dropout (Srivastava et al., 2014) and a baseline method using MAP training.

2 Training for Robustness

Our method introduces a regularization term during training to ensure model robustness. We develop our approach based on a general class of parametric mod-

¹Implementation available at <https://github.com/lrank/Robust-Representation>.

els, with the following structure. Let \mathbf{x} be the input, which is a sequence of (discrete) words, represented by a fixed-size vector of continuous values, \mathbf{h} . A transfer function takes \mathbf{h} as input and produces an output distribution, \mathbf{y}_{pred} . Training proceeds using stochastic gradient descent to minimize a loss function L , measuring the difference between \mathbf{y}_{pred} and the truth \mathbf{y}_{true} .

The purpose of our work is to learn neural models which are more robust to strange or invalid inputs. When small perturbations are applied on \mathbf{x} , we want the prediction \mathbf{y}_{pred} to remain stable. Text can be highly variable, allowing for the same information to be conveyed with different word choice, different syntactic structures, typographical errors, stylistic changes, etc. This is a particular problem in transfer learning scenarios such as domain adaptation, where the inputs in distinct domains are drawn from related, but different, distributions. A good model should be robust to these kinds of small changes to the input, and produce reliable and stable predictions.

Next we discuss methods for learning models which are robust to variations in the input, before providing details of the neural network model used in our experimental evaluation.

2.1 Conventional Regularization and Dropout

Conventional methods for learning robust models include l_1 and l_2 regularization (Ng, 2004), and dropout (Srivastava et al., 2014). In fact, Wager et al. (2013) showed that the dropout regularizer is first-order equivalent to an l_2 regularizer applied after scaling the features. Dropout is also equivalent to “Follow the Perturbed Leader” (FPL) which perturbs exponential numbers of experts by noise and then predicts with the expert of minimum perturbed loss for online learning robustness (van Erven et al., 2014). Given its popularity in deep learning, we take dropout to be a strong baseline in our evaluation.

The key idea behind dropout is to randomly zero out units, along with their connections, from the network during training, thus limiting the extent of co-adaptation between units. We apply dropout on the representation vector \mathbf{h} , denoted $\hat{\mathbf{h}} = \text{dropout}_{\beta}(\mathbf{h})$, where β is the dropout rate. Similarly to our proposed method, training with dropout requires gradient based search for the minimizer of the loss L .

We also use dropout to generate noise in the test

data as part of our experimental simulations, as we will discuss later.

2.2 Robust Regularization

Our method is inspired by the work on adversarial training in computer vision (Goodfellow et al., 2014). In image recognition tasks, small distortions that are indiscernible to humans can significantly distort the predictions of neural networks (Szegedy et al., 2014). An intuitive explanation of our regularization method is, when noise is applied to the data, the variation of the output is kept lower than the noise. We adapt this idea from Rifai et al. (2011) and develop the Jacobian regularization method.

The proposed regularization method works as follows. Conventional training seeks to minimise the difference between \mathbf{y}_{true} and \mathbf{y}_{pred} . However, in order to make our model robust against noise, we also want to minimize the variation of the output when noise is applied to the input. This is to say, when perturbations are applied to the input, there should be as little perturbation in the output as possible. Formally, the perturbations of output can be written as $\mathbf{p}_y = M(\mathbf{x} + \mathbf{p}_x) - M(\mathbf{x})$, where \mathbf{x} is the input, \mathbf{p}_x is the vector of perturbations applied to \mathbf{x} , M expresses the trained model, \mathbf{p}_y is the vector of perturbations generated by the model, and the output distribution $\mathbf{y} = M(\mathbf{x})$. Therefore

$$\lim_{\mathbf{p}_x \rightarrow \mathbf{0}} \mathbf{p}_y = \lim_{\mathbf{p}_x \rightarrow \mathbf{0}} (M(\mathbf{x} + \mathbf{p}_x) - M(\mathbf{x})) = \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \cdot \mathbf{p}_x,$$

$$\text{and distance} \left(\lim_{\mathbf{p}_x \rightarrow \mathbf{0}} \mathbf{p}_y / \mathbf{p}_x, \mathbf{0} \right) = \left\| \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \right\|_F.$$

In other words, minimising local noise sensitivity is equivalent to minimising the Frobenius norm of the Jacobean matrix of partial derivatives of the model outputs *wrt* its inputs.

To minimize the effect of perturbation noise, our method involves an additional term in the loss function, in the form of the derivative of loss L with respect to hidden layer \mathbf{h} . Note that while in principle we could consider robustness to perturbations in the input \mathbf{x} , the discrete nature of \mathbf{x} adds additional mathematical complications, and thus we defer this setting for future work. Combining the elements, the new loss function can be expressed as

$$\mathcal{L} = L + \lambda \cdot \left\| \frac{\partial L}{\partial \mathbf{h}} \right\|_2, \quad (1)$$

where λ is a weight term, and distance takes the form of the l_2 norm. The training objective in Equation (1) supports gradient optimization, but note that it requires the calculation of second-order derivatives of L during back propagation, arising from the $\partial L/\partial \mathbf{h}$ term. Henceforth we refer to this method as **robust regularization**.

2.3 Convolutional Network

For the purposes of this paper, we focus exclusively on convolutional neural networks (CNNs), but stress that the method is compatible with other neural architectures and other types of parametric models (not just deep neural networks). The CNN used in this research is based on the model proposed by Kim (2014), and is outlined below.

Let S be the sentence, consisting of n words $\{w_1, w_2, \dots, w_n\}$. A look-up table is applied to S , made up of word vectors $\mathbf{e}_i \in \mathbb{R}^m$ corresponding to each word w_i , where m is the word vector dimensionality. Thus, sentence S can be represented as a matrix $\mathbf{E}_S \in \mathbb{R}^{m \times n}$ by concatenating the word vectors $\mathbf{E}_S = \bigoplus_{i=1}^n \mathbf{e}_{w_i}$.

A convolutional layer combined with a number of wide convolutional filters is applied to \mathbf{E}_S . Specifically, the k -th convolutional filter operator filter_k involves a weight vector $\mathbf{w}_k \in \mathbb{R}^{m \times t}$, which works on every t_k -sized window of \mathbf{E}_S , and is accompanied by a bias term $b \in \mathbb{R}$. The filter operator is followed by the non-linear function \mathcal{F} , a rectified linear unit, ReLU, followed by a max-pooling operation, to generate a hidden activation $h_k = \text{MaxPooling}(\mathcal{F}(\text{filter}_k(\mathbf{E}_S; \mathbf{w}_k, b)))$. Multiple filters with different window sizes are used to learn different local properties of the sentence. We concatenate all the hidden activations h_k to form a hidden layer \mathbf{h} , with size equal to the number of filters. Details of parameter settings can be found in Section 3.2.

The feature vector \mathbf{h} is fed into a final softmax layer with a linear transform to generate a probability distribution over labels

$$\mathbf{y}_{\text{pred}} = \text{softmax}(\mathbf{w} \cdot \mathbf{h} + \mathbf{b}),$$

where \mathbf{w} and \mathbf{b} are parameters. Finally, the model minimizes the loss of the cross-entropy between the ground-truth and the model prediction,

$L = \text{CrossEntropy}(\mathbf{y}_{\text{true}}, \mathbf{y}_{\text{pred}})$, for which we use stochastic gradient descent.

3 Datasets and Experimental Setups

We experiment on the following datasets,² following Kim (2014):

- MR: Sentence polarity dataset (Pang and Lee, 2008)³
- Subj: Subjectivity dataset (Pang and Lee, 2005)³
- CR: Customer review dataset (Hu and Liu, 2004)⁴
- SST: Stanford Sentiment Treebank, using the 3-class configuration (Socher et al., 2013)⁵

In each case, we evaluate using classification accuracy.

3.1 Noisifying the Data

Different to conventional evaluation, we corrupt the test data with noise in order to evaluate the robustness of our model. We assume that when dealing with short text such as Twitter posts, it is common to see unknown words due to typos, abbreviations and sociolinguistic marking of different types (Han and Baldwin, 2011; Eisenstein, 2013). To simulate this, we apply word-level dropout noise to each document, by randomly replacing words by a unique sentinel symbol.⁶ This is applied to each word with probability $\alpha \in \{0, 0.1, 0.2, 0.3\}$.

We also experimented with adding different levels of Gaussian noise to the sentence embeddings \mathbf{E}_S , but found the results to be largely consistent with those for word dropout noise, and therefore we have omitted these results from the paper.

To directly test the robustness under a more realistic setting, we additionally perform cross-domain evaluation, where we train a model on one dataset

²For datasets where there is no pre-defined training/test split, we evaluate using 10-fold cross validation. Refer to Kim (2014) for more details on the datasets.

³<https://www.cs.cornell.edu/people/pabo/movie-review-data/>

⁴<http://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html>

⁵<http://nlp.stanford.edu/sentiment/>

⁶This was to avoid creating new n -grams which would occur when symbols are deleted from the input. Masking tokens instead results in partially masked n -grams as input to the convolutional filters.

Dataset		MR				Subj			
Word dropout rate (α)		0	0.1	0.2	0.3	0	0.1	0.2	0.3
Baseline		80.5	79.4	77.9	76.5	93.1	92.0	90.9	89.8
Dropout (β)	0.3	80.3	79.5	78.1	76.7	92.7	92.0	90.9	89.5
	0.5	80.3	79.0	78.0	76.5	93.0	92.0	91.1	89.9
	0.7	80.3	79.3	78.3	76.8	92.8	91.9	90.9	89.8
Robust Regularization (λ)	10^{-3}	80.5	79.4	78.3	76.7	93.0	92.2	91.1	89.8
	10^{-2}	80.8	79.3	78.4	77.0	93.0	92.2	91.0	90.0
	10^{-1}	80.4	78.8	77.8	77.0	92.7	91.9	91.0	89.8
	1	79.3	77.1	76.1	75.5	91.7	91.1	90.1	89.3
Dropout + Robust	$\beta = 0.5, \lambda = 10^{-2}$	80.6	79.9	78.6	77.3	93.0	92.2	91.2	90.1

Dataset		CR				SST			
Word dropout rate (α)		0	0.1	0.2	0.3	0	0.1	0.2	0.3
Baseline		83.2	82.3	80.4	77.9	84.1	82.3	80.3	77.8
Dropout (β)	0.3	83.3	82.1	80.3	78.9	84.2	82.3	80.2	78.0
	0.5	83.2	82.4	81.0	79.3	84.2	82.4	80.5	78.2
	0.7	83.2	82.2	80.7	78.8	83.9	82.5	80.9	78.2
Robust Regularization (λ)	10^{-3}	83.3	82.6	81.4	79.5	84.5	82.8	81.4	78.8
	10^{-2}	83.4	82.5	81.6	79.3	84.2	82.4	80.7	78.6
	10^{-1}	83.3	82.7	82.0	79.6	82.5	81.5	79.7	77.6
	1	82.9	81.4	79.8	79.0	82.2	80.9	79.1	77.3
Dropout + Robust	$\beta = 0.5, \lambda = 10^{-2}$	83.3	82.5	81.5	79.7	84.3	82.6	80.8	79.1

Table 1: Accuracy (%) with increasing word-level dropout across the four datasets. For each dataset, we apply four levels of noise $\alpha = \{0, 0.1, 0.2, 0.3\}$; the best result for each combination of α and dataset is indicated in **bold**. The Baseline model is a simple CNN model without regularization. The last model combines dropout and our method with fixed parameters β and λ as indicated.

and apply it to another. For this, we use the pairing of MR and CR, where the first dataset is based on movie reviews and the second on product reviews, but both use the same label set. Note that there is a significant domain shift between these corpora, due to the very nature of the items reviewed.

3.2 Word Vectors and Hyper-parameters

To set the hyper-parameters of the CNN, we follow the guidelines of Zhang and Wallace (2015), setting word embeddings to $m = 300$ dimensions and initialising based on word2vec pre-training (Mikolov et al., 2013). Words not in the pre-trained vector table were initialized randomly by the uniform distribution $U([-0.25, 0.25]^m)$. The window sizes of filters (t) are set to 3, 4, 5, with 128 filters for each size, resulting in a hidden layer dimensionality of $384 = 128 \times 3$. We use the Adam optimizer (Kingma and Ba, 2015) for training.

4 Results and Discussions

The results for word-level dropout noise are presented in Table 1. In general, increasing the word-level dropout noise leads to a drop in accuracy for all four datasets, however the relative dropoff in accuracy for Robust Regularization is less than for Word Dropout, and in 15 out of 16 cases (four noise levels across the four datasets), our method achieves the best result. Note that this includes the case of $\alpha = 0$, where the test data is left in its original form, which shows that Robust Regularization is also an effective means of preventing overfitting in the model.

For each dataset, we also evaluated based on the combination of Word Dropout and Robust Regularization using the fixed parameters $\beta = 0.5$ and $\lambda = 10^{-2}$, which are overall the best individual settings. The combined approach performs better than either individual method for the highest noise levels tested across all datasets. This indicates that Robust

Train/Test		MR/CR	CR/MR
Baseline		67.5	61.0
Dropout (β)	0.3	71.6	62.2
	0.5	71.0	62.1
	0.7	70.9	62.0
Robust Regularization (λ)	10^{-3}	70.8	61.6
	10^{-2}	71.1	62.5
	10^{-1}	72.0	62.2
	1	71.8	62.3
Dropout + Robust	$\beta = 0.5, \lambda = 10^{-2}$	72.0	62.4

Table 2: Accuracy under cross-domain evaluation; the best result for each dataset is indicated in **bold**.

Regularization acts in a complementary way to Word Dropout.

Table 2 presents the results of the cross-domain experiment, whereby we train a model on MR and test on CR, and vice versa, to measure the robustness of the different regularization methods in a more real-world setting. Once again, we see that our regularization method is superior to word-level dropout and the baseline CNN, and the techniques combined do very well, consistent with our findings for synthetic noise.

4.1 Running Time

Our method requires second-order derivatives, and thus is a little slower at training time. Figure 1 is a plot of the training and test accuracy at varying points during training over SST.

We can see that the runtime till convergence is only slightly slower for Robust Regularization than standard training, at roughly 30 minutes on a two-core CPU (one fold) with standard training vs. 35–40 minutes with Robust Regularization. The convergence time for Robust Regularization is comparable to that for Word Dropout.

5 Conclusions

In this paper, we present a robust regularization method which explicitly minimises a neural model’s sensitivity to small changes in its hidden representation. Based on evaluation over four sentiment analysis datasets using convolutional neural networks, we found our method to be both superior and complementary to conventional word-level dropout under varying levels of noise, and in a cross-domain evalu-

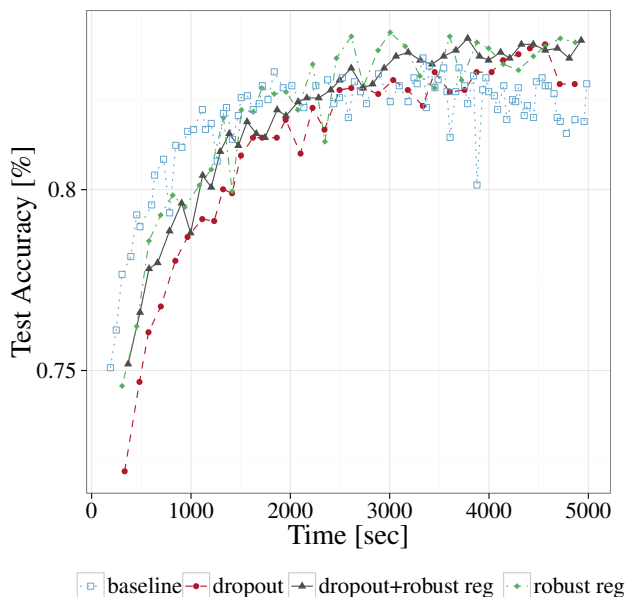


Figure 1: Time–accuracy evaluation over the different combinations of Word Dropout (dropout) and Robust Regularization (robust reg) over SST, without injecting noise.

ation.

For future work, we plan to apply our regularization method to other models and tasks to determine how generally applicable our method is. Also, we will explore methods for more realistic linguistic noise, such as lexical, syntactic and semantic noise, to develop models that are robust to the kinds of data often encountered at test time.

Acknowledgments

We are grateful to the anonymous reviewers for their helpful feedback and suggestions. This work was supported by the Australian Research Council (grant numbers FT130101105 and FT120100658). Also, we would like to thank the developers of Tensorflow (Abadi et al., 2015), which was used for the experiments in this paper.

References

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg,

- Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-scale machine learning on heterogeneous systems. Technical report, Google Research.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the International Conference on Learning Representations*.
- Zsolt Bitvai and Trevor Cohn. 2015. Non-linear text regression with a deep convolutional neural network. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Short Papers)*, pages 180–185.
- Jacob Eisenstein. 2013. What to do about bad language on the internet. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 359–369.
- Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2015. Analysis of classifiers’ robustness to adversarial perturbations. *arXiv preprint arXiv:1502.02590*.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. In *Proceedings of the International Conference on Learning Representations*.
- Alan Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6645–6649.
- Shixiang Gu and Luca Rigazio. 2014. Towards deep neural network architectures robust to adversarial examples. In *Proceedings of the NIPS 2014 Deep Learning and Representation Learning Workshop*.
- Bo Han and Timothy Baldwin. 2011. Lexical normalization of short text messages: Mkn sens a #twitter. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 368–378.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 168–177.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 655–665.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1746–1751.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*, pages 1097–1105.
- James Martens. 2010. Deep learning via Hessian-free optimization. In *Proceedings of the 27th International Conference on Machine Learning*, pages 735–742.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26*, pages 3111–3119.
- Andrew Y. Ng. 2004. Feature selection, L1 vs. L2 regularization, and rotational invariance. In *Proceedings of the Twenty-first International Conference on Machine Learning*.
- Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 115–124.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135.
- Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. 2011. Contractive auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the 28th International Conference on Machine Learning*, pages 833–840.
- Richard Socher, Alex Perelygin, Jean Y. Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob

- Fergus. 2014. Intriguing properties of neural networks. In *Proceedings of the International Conference on Learning Representations*.
- Pedro Tabacof and Eduardo Valle. 2016. Exploring the space of adversarial images. In *Proceedings of the IEEE International Joint Conference on Neural Networks*.
- Tim van Erven, Wojciech Kotłowski, and Manfred K. Warmuth. 2014. Follow the leader with dropout perturbations. In *Proceedings of the 27th Conference on Learning Theory*, pages 949–974.
- Stefan Wager, Sida Wang, and Percy S. Liang. 2013. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems 26*, pages 351–359.
- Wen-tau Yih, Xiaodong He, and Christopher Meek. 2014. Semantic parsing for single-relation question answering. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Short Papers)*, pages 643–648.
- Ye Zhang and Byron Wallace. 2015. A sensitivity analysis of (and practitioners’ guide to) convolutional neural networks for sentence classification. *arXiv preprint arXiv:1510.03820*.