
Dropout as data augmentation

Kishore Konda

Goethe University Frankfurt, Germany
konda.kishorereddy@gmail.com

Xavier Bouthillier

Université de Montréal, Canada
xavier.bouthillier@umontreal.ca

Roland Memisevic

Université de Montréal, Canada
roland.memisevic@umontreal.ca

Pascal Vincent

Université de Montréal, Canada and CIFAR
pascal.vincent@umontreal.ca

Abstract

We show that using dropout in a network can be interpreted as a kind of data augmentation without domain knowledge. We present an approach to projecting the dropout noise within a network back into the input space, thereby generating augmented versions of the training data, and we show that training a deterministic network on the augmented samples yields similar results. We furthermore propose an explanation for the increased sparsity levels that can be observed in networks trained with dropout and show how this is related to data augmentation. Finally, we detail a random dropout noise scheme based on our observations and show that it improves dropout results without adding significant computational cost.

1 Introduction

Noise is normally seen as intrinsically undesirable. The word itself bears a very negative connotation. It is not surprising then that many early mathematical models in neuroscience aimed to factor out noise by any means. A few decades ago, the use of stochastic resonance [1] in neuro-scientific models initiated a new interest in neuroscience regarding random fluctuations and the role they play in the brain. Theories about neuronal noise are now flourishing and previous deterministic models are improved by the incorporation of noise [2].

Biological brains have always been a strong inspiration when it comes to developing learning algorithms. Considering the amount of noise which takes place in the brain during learning, one can wonder if this has any beneficial effect. Many techniques in machine learning have made use of noise to improve performance recently, namely, Denoising Autoencoders [3], dropout [4] and its relative, DropConnect [5]. Those successful approaches suggest that neuronal noise plays a fundamental role in the process of learning and should be studied more thoroughly.

Using dropout can be viewed as training a huge number of neural networks with shared parameters and applying bagging at test time for better generalization [6]. Binary noise can also be viewed as preventing neurons from co-adapting, which improves the generalization of the model even more. In this paper, we propose an alternative view inspired by results coming from computational neuroscience. Yarom and Hounsgaard [2] reports that the variation in population of neuronal activity is believed to express ambiguity about the environment¹. Such variations can also be thought of as a basic form of data augmentation without domain knowledge. We propose that multiplicative noise schemes like dropout are implicitly incorporating a form of sophisticated data augmentation. In Section 3, we formulate a method to generate data which replicates dropout noise within a deterministic network, and we demonstrate in Section 5 that there is no significant loss of accuracy.

¹The environment is basically the input of the network

Recent work [7] has shown that regularization in autoencoders, also including dropout style input corruption, leads to tiling of the input space² via strong negative hidden biases. In the present work, the idea is extended to multilayer discriminative networks with the hypothesis that input space tiling is the best possible representation when the input samples are distributed along a highly non-linear or discontinuous manifold. We also show that using dropout noise, which acts like data augmentation, yields in a more fine-grained tiling of the input space resulting in sparser representations and better classification performance. We experimentally show that the magnitude of the sparsity of the hidden representation is proportional to the corresponding input noise level. An explanation for this observation is presented in Section 4.

Finally, capitalizing on the idea of data augmentation, we present an extension of dropout which uses random noise levels to improve the variety of samples. This simple extension gives better classification performance across different network architectures, yielding competitive results on the MNIST permutation invariant classification task.

2 Dropout and its derivatives

Dropout [4] is a technique where a fraction of hidden units in a deep network are randomly turned off while training. Given an input vector \vec{x} to a layer and an N -dimensional weight vector \vec{w}_n the hidden response corresponding to filter \vec{w}_n before the non-linearity can be written as

$$\vec{h}_n = \sum_i x_i w_{n,i} + b = \vec{w}_n^T \vec{x} + b_n \quad (1)$$

Using dropout involves multiplying \vec{x} with a binary mask sampled with probability p . Most often, the probability used is 0.5 which implies half of the input dimensions (or hidden units if the input comes from the previous layer) are dropped. To compensate for the loss of magnitude the active dimensions are scaled by $\frac{1}{p} = 2$. The new hidden value can be written

$$h_n = \vec{w}_n^T (\vec{r} \odot \vec{x}) + b_n \quad (2)$$

where $P(r_i = 1) = 0.5$. Similarly, DropConnect can be formalized as

$$h_n = (\vec{r}_n \odot \vec{w}_n)^T \vec{x} + b_n \quad (3)$$

3 Dropout from a data augmentation perspective

In many previous works it has been shown that augmenting data by using domain specific transformations helps in learning better models [8, 9, 10, 11]. In this work, we analyze dropout in the context of data augmentation. Considering the task of classification, given a set of training samples, the objective would be to learn a mapping function which maps every input to its corresponding output label. To generalize, the mapping function needs to be able to correctly map not just the training samples but also any other samples drawn from the data distribution. This means that it must not only map input space sub-regions represented by training samples, but all sub-regions containing a significant amount of data. One way to learn such a mapping function is by augmenting the training data such that it covers a larger portion of the natural distribution. Domain-based data augmentation helps to artificially boost training data coverage which makes it possible to train a better mapping function. We hypothesize that noise based regularization techniques result in a similar effect of increasing training data coverage at every hidden layer and this work presents multiple experimental observations to support our hypothesis.

3.1 Projecting noise back into the input space

An intuitive view of dropout is that the noise at each hidden layer is a random transformation of the input at different levels of representation. It is less clear however how those transformations affect the network globally.

²By this, we mean that the features learned by the layer will cover small non-overlapping sub-regions of the input space, thus dividing it into tiles.

Suppose that for each transformation, it would be possible to find an input which would approximately generate the same hidden activity deterministically. This would make it possible to remove the noise from the network and to train it deterministically on transformed data, by projecting the noise back into the input.

Assume $h(x)$ is a linear projection of a d_i -dimensional input x into a d_h -dimensional space:

$$h(x) = xW + b \quad (4)$$

and given $a(h)$ and $\tilde{a}(h)$, an activation function and its noisy version where \vec{m} is a random variable sampled from a Bernoulli distribution $B(p_h)$ and $\text{rect}(h)$, a rectifier

$$a(h) = \text{rect}(h) \quad (5)$$

$$\tilde{a}(h) = \vec{m} \odot \text{rect}(h) \quad (6)$$

We assume that for a given $\tilde{a}(h)$, there exist an x^* , **such that**

$$(a \circ h)(x^*) = \text{rect}(h(x^*)) \approx \vec{m} \odot \text{rect}(h(x)) = (\tilde{a} \circ h)(x) \quad (7)$$

Similarly to adversarial examples from [12], an x^* can be found by minimizing the squared error L using stochastic gradient descent

$$L(x, x^*) = |(a \circ h)(x^*) - (\tilde{a} \circ h)(x)|^2 \quad (8)$$

Equation 8 can be generalized to a network with n hidden layers. To lighten notation we first define

$$\tilde{f}^{(i)}(x) = \left(\tilde{a}^{(i)} \circ h^{(i)} \circ \dots \circ \tilde{a}^{(1)} \circ h^{(1)} \right) (x) \quad (9)$$

$$f^{(i)}(x^*) = \left(a^{(i)} \circ h^{(i)} \circ \dots \circ a^{(1)} \circ h^{(1)} \right) (x^*) \quad (10)$$

which results in generalized equation

$$L(x, x^{(1)*}, \dots, x^{(n)*}) = \sum_{i=1}^n \lambda_i \left| f^{(i)}(x^{(i)*}) - \tilde{f}^{(i)}(x) \right|^2 \quad (11)$$

We can show, with a proof by contradiction, that it's **unlikely** to find a single $x^* = x^{(1)*} = x^{(2)*} = \dots = x^{(n)*}$ that minimizes well L . Proof is detailed in appendix subsection 7.1. Fortunately, it is easy to find a different x^* for each hidden layer, by providing multiple inputs $(x, x^{(1)*}, x^{(2)*}, \dots, x^{(n)*})$, where n is the number of hidden layers. As each $x^{(i)*}$ is the projection of a transformation in the representation space defined by the i -th hidden layer, it suggests viewing dropout as a sophisticated data augmentation procedure that **samples** data around training examples with respect to different level of representations. This raises the question whether we can train the network deterministically on the $x^{(i)*}$ rather than using dropout. The answer is not trivial, because

1. As training is now deterministic, the stimuli will spread inside the network without any noise distortion, making the hypothesis that dropout avoids co-adaptation less likely.
2. The gradients of the linear projections will differ greatly. In the case of dropout, $\frac{\partial h}{\partial W^{(i)}}$ is always equal to its input transformation, i.e. $\tilde{f}^{(i-1)}(x)$, whereas the deterministic version of the training will update $W^{(i)}$ according to $(f^{(i-1)}(x^{(1)*}), \dots, f^{(i-1)}(x^{(n)*}))^3$.

Although we proved a single x^* minimising 11 is difficult to find for a large network, we show experimentally in section 5 that it is possible to do so within reasonable approximation for a relatively small two hidden layers network. We further show that dropout can be **replicated** by projecting the noise back on the input space without a significant loss of accuracy.

³Because we train on n samples from x , one for each hidden layer

4 Sparsity and input space tiling

We explained above how training a network with dropout at different layers can also be viewed approximately as augmenting the training data, but in a smart way that depends on the network parameters and chosen level of representation. In this section we present additional support for this argument based on experimental observations about the sparsity of the hidden layer representations in a network. As we will see in the experiments section 5, in a multi-layered network trained with dropout, it can be observed that the sparsity of the hidden layer representation, given sufficient hidden units, is roughly proportional to the corresponding input corruption level. In the section we present a hypothesis to explain this behavior.

4.1 Non-linear manifolds and tiling

When the data points belonging to a particular class are distributed along a linear manifold, or sub-space, of the input space, it is enough to learn a single set of features which can span the entire manifold. But when the data is distributed along a highly non-linear and discontinuous manifold, the best way to represent such a distribution is to learn features which can explicitly represent small local regions of the input space, effectively “tiling” the space to define non-linear decision boundaries.

4.2 Why sparsity?

Let us assume that an input distribution is represented by M different discontinuous manifolds or clusters in an N dimensional input space. Given a set of K feature vectors an efficient representation would be one which maps all input clusters to representations with least overlap (in terms of hidden dimensions) across multiple clusters. Let A_m be the set of active dimensions of the hidden representation that represents a given sub-space m . For a sufficiently large K , even with large A_m , one could generate enough hidden representations to map all the clusters such that the overlap across representations $A_i \cap A_j; i, j \in \{1, \dots, M\}$ is small. This would correspond to a dense representation as the set of active dimensions A_m is large. But when we have a small K and large M , which is mostly the case in practice, the only way to generate hidden representations with small overlap is to lower the active subset size A_m i.e., increase the sparsity of the representation. When K is insufficient to generate representations with least overlap it results in denser and inefficient representations with larger overlap.

To illustrate this view we use a classification experiment on a 2D toy-dataset. We define a set of data points distributed on unit circle with different small arcs on the circle belonging to different classes. For simplicity all the arcs are of same length. Any two arcs of the same class are separated by arcs of other classes creating discontinuous manifolds.

We defined multiple datasets by varying the arc length such that in each dataset the number of input clusters is different. The datasets are defined such that the smaller the length of the arc, the higher the number of arcs per class overall. The circumference of the circle is defined by 15000 points and the total number of classes is 10. Two example datasets are shown in Figure 1. We train networks with a single hidden layer containing 12 units to perform classification. The training data in the case of each dataset is a set of 5000 points randomly selected from the total set of points. We also ran the same experiments using training set size of 1000 points. The sparsity of the hidden layer representations from the networks trained on different datasets is plotted in Figure 1. From the plot it can be observed that the sparsity of the hidden representation decreases with increase in arc-length, i.e., with decreasing number of clusters the model needs to map. It can also be observed that the sparsity is not proportional to the number of training samples but mainly to the number of clusters.

Since higher dropout levels also result in sparser representations (see Section 5), we hypothesize that the purpose of corruption can also be viewed as creating more input sub-spaces or clusters to cover by effectively augmenting the original data samples. This is in line with the data augmentation argument from the previous section.

4.3 Random noise levels

In almost all previous work using corruption as regularization, the level of corruption for each layer is fixed during training. Some recent work reports that noise level schedules, inspired by simulated




Figure 1: (a), (b); Two sample input distributions defined with different arc lengths. (c) The sparsity of the hidden layer from a single layer MLP trained on datasets defined using different arc lengths.

annealing, help in supervised and unsupervised tasks [13, 14, 15]. Assuming that corruption implies data augmentation, a **variable** noise level for each layer will result in a richer augmentation of the training data, i.e., more diverse transformations. While this is not the objective of a decreasing schedule, it does have a diversification effect on the transformations.

We propose an alternative that avoids a schedule and rather uses random noise levels such that the model cannot adapt to slowly changing noise distribution.

Random-dropout At each layer the corruption level p is sampled per training sample from a uniform distribution in the interval $[a, b]$, $0 \leq a < b \leq 1$

A similar approach of sampling the noise level was used in [13] in the context of unsupervised learning using an autoencoder. However, they show that the approach is not very useful in their case.

4.4 Thresholding non-linearities

If learned features “tile” the input space, one would expect that the information whether a hidden unit is on or off is more important than the actual magnitude of the hidden response. To test this hypothesis we performed additional experiments using a thresholding non-linearity, *rectified-tanh*. The rectified-tanh is obtained by applying rectification following the tanh function which gives the positive half of the tanh function. We ran multiple experiments with different noise schemes using the same approaches that were used for ReLU networks in Section 5. A network with 2500-1250 hidden units trained with dropout noise (0.2 input noise, 0.5 hidden noise) gave 1.02% error on MNIST dataset and 39.61% on CIFAR-10. The same experiment using Random-Dropout with noise interval of $[0.0, 0.5]$ gave 0.91% error on MNIST and 38.42% on CIFAR-10.

An interesting fact that emerged from our experiments is that models with ReLU activation functions are on-par with models using rectified-tanh. As the latter function quickly saturates to 1, it seems that the linear part of the ReLU is not fully exploited. We made an experiment to test if the hypothesis is true. We trained a rectified-tanh network with 2500-1250 hidden units with dropout noise (0.2 input noise, 0.5 hidden noise) on MNIST and CIFAR-10 datasets. For inference, the response h (after non-linearity) of each hidden unit is set to 1 if $h > 0$ and 0 when $h \leq 0$, turning the unit binary. The binary inference resulted in 1.41% test error on MNIST and 41.02% on CIFAR-10. This shows that the network relies more on a unit being on or off than the magnitude of its response. Moreover, this is consistent with the hypothesis that a sparse representation is tiling input sub-regions, i.e. the activation represents whether an example is within a given sub-region or not.

5 Experiments

We ran a series of experiments with Multi-Layer-Perceptron (MLP) architectures on the MNIST and CIFAR-10 datasets to study the effect of noise schemes on the performance and the learned network parameters. Each MLP consists of two hidden layers with ReLU units followed by a *softmax* layer.




Figure 2: Error percentages on the MNIST and CIFAR-10 datasets using MLP architectures trained with different corruption schemes. **Row-1:** Experiments on CIFAR-10. **Row-2:** Experiments on MNIST. **Column-1:** Using dropout with varying input noise and fixed hidden noise of 0.5. **Column-2:** Using dropout with varying hidden noise with fixed input noise of 0.2. **Column-3:** Using Random-dropout with varying noise range $[0, x]$ used at hidden and input layers.

We experimented with three different MLP architectures each one with a different number of units in the hidden layers: 2500-625, 2500-1250 and 2500-2500. The MNIST dataset consists of 60000 training samples and 10000 test samples. We split the training set into a set of 50000 samples for training and 10000 for validation. Each network is trained for 501 epochs and the best model based on validation error is selected. Later the best model is further trained on the complete training set of 60000 samples (training + validation split) for another 150 epochs. The mean error percentage for the last 75 epochs is used for comparison.

We also ran experiments on the CIFAR-10 permutation invariant task using the same MLPs described above. The dataset consists of 50000 training samples and 10000 test samples. We use PCA based dimensionality reduction without whitening as preprocessing scheme. We used the same approach as that in the MNIST experiments to train the networks and for reporting the performances.

First, we evaluated the dropout noise scheme by training the MLPs with a fixed hidden noise level of 0.5 and the input noise level varying from 0.0 to 0.7 with increments of 0.1 for each experiment. In the second experiment, we fixed the input noise level at 0.2 and the hidden noise level is varied from 0.0 to 0.7, again with an increment of 0.1. In the final set of experiments we use the Random Dropout noise scheme using the same noise level at input and hidden layers. The noise level in this case is a range $[0, x]$ where x is varied from 0.0 to 0.8 with increment 0.1. The classification performances corresponding to the all the experiments on both the datasets are reported in Figure 2.

The plots showing the **sparsity** of hidden layer representations against **noise** levels are shown in Figure 3. In that figure only plots corresponding to the random noise scheme experiments are shown where it can be observed that **the sparsity of a given hidden layer representation is proportional to the level of noise applied to its input.** We observed the same behavior in case of the dropout experiments. To show that the proportionality between noise level and sparsity is not limited to using dropout noise, we performed an additional experiment using Gaussian noise for corruption. We trained a two-layered MLP with 2500 – 1250 hidden units on MNIST dataset using zero-mean additive Gaussian noise for input corruption. The standard deviation of the Gaussian is varied to be $[0., 0.2, .0.4, 0.6]$. The corresponding layer-1 sparsity levels are shown in figure 3.

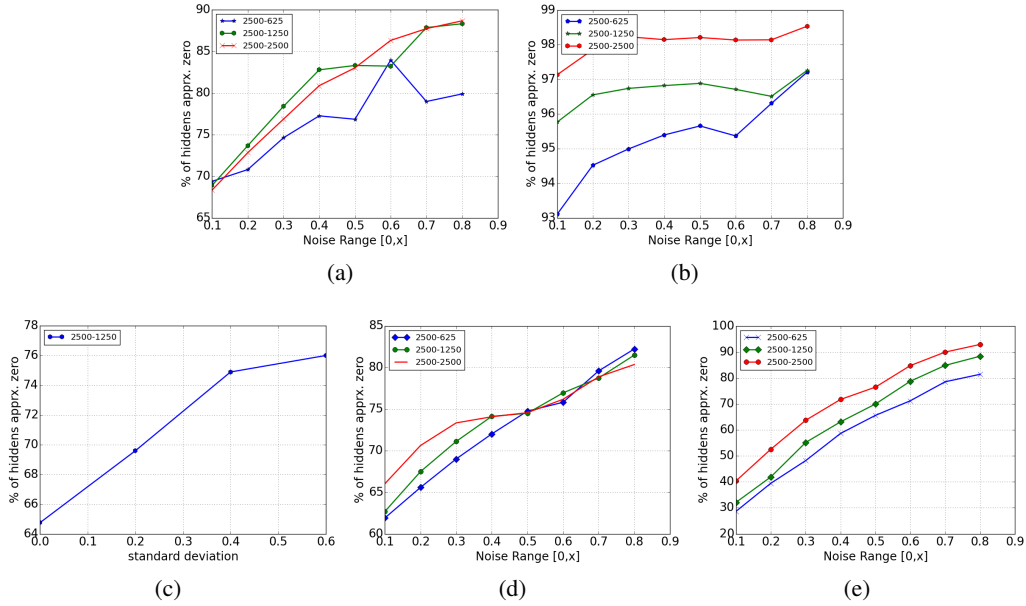


Figure 3: Sparsity in different layers of MLPs trained using noise based regularization schemes. **(a)** and **(b)** correspond to sparsity in layer-1 and layer-2 respectively of networks trained on CIFAR-10 using random-dropout. **(d)** and **(e)** show the corresponding results for MNIST. **(c)** shows the layer-1 sparsity of a network trained using zero-mean additive Gaussian noise for input corruption on MNIST.

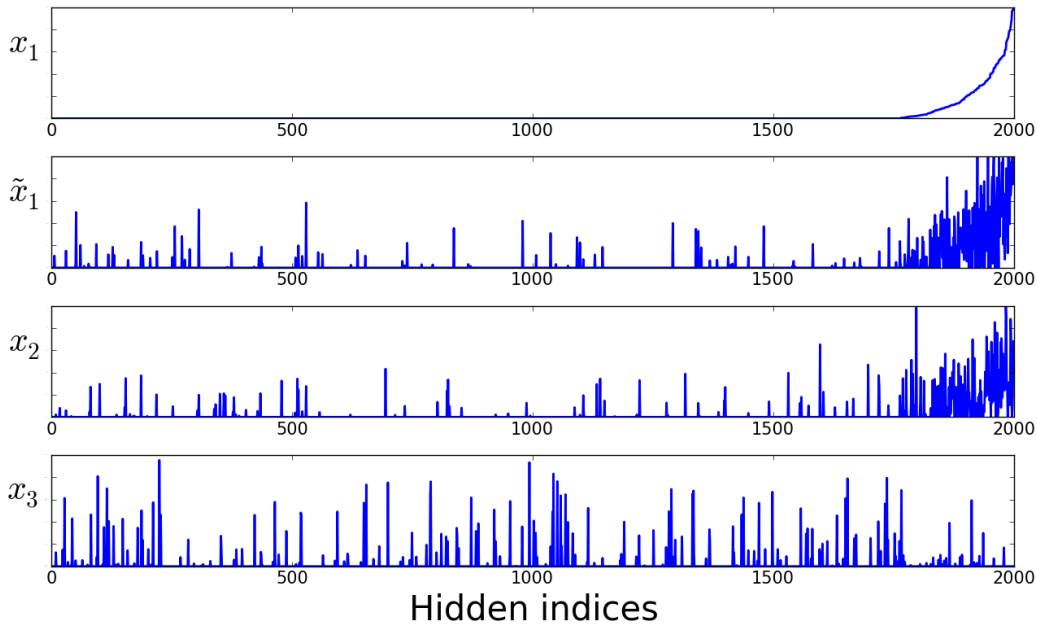


Figure 4: Comparing layer-2 responses of a 1000-2000-3000 unit network trained with dropout corruption on CIFAR-10. x_1 : Original train sample without dropout. \tilde{x}_1 : x_1 when using regular dropout in the network. x_2 : Different sample from same class as that of x_1 . x_3 : Sample from a different class. For x_1 , the indices of hidden representation are sorted in increasing order of the corresponding response and the same order is used for hidden representations in other cases.

EXPERIMENTAL SETUP	MNIST	CIFAR-10
DROPOUT ON INPUT	1.12%	40.9%
DROPOUT ON INPUT+HIDDEN	0.95%	39.0%
NOISE PROJECTION	0.99%	37.9%

Table 1: Test accuracy on permutation invariant MNIST and CIFAR-10 (2500-1250 hidden units with 0.2 dropout input noise and 0.5 hidden layer dropout noise where applicable). Noise on CIFAR-10 is applied to the PCA representation.

5.1 Training deterministically with noise projected back into input

We used a similar setting to test the effect of replacing dropout by **back-projected inputs**. We trained MLPs with 2500 hidden units in the first and 1250 hidden units in the second layer in all experiments. We used input noise level 0.2 and dropout level 0.5 for both hidden layers in our reference models and the same values to compute back-projected inputs.

At each epoch, an x^* is generated for each training sample. It proved to be possible to find good x^* approximations for the entire network at once for a 2-hidden layer network. Thus, we trained on x and x^* solely rather than x , $x^{(1)*}$ and $x^{(2)*}$ as it gave a significant speed up. For simplicity, the network is trained on x for an epoch than on x^* for an epoch. All x^* are generated with parameter values of the model at the beginning of the epoch.

Noisy inputs x^* are found using stochastic gradient descent. 20 learning steps are done with a learning rate of 300.0 for first hidden layer and 30 for second hidden layer.

In order to get comparable results for input transformations, we also report results with Bernoulli binary mask applied on the input. Results for simple input transformations, dropout and noise projection are shown in table 1.

Results from table 1 suggest that Dropout can be replicated by projecting the noise back into the input and training a neural network deterministically on this generated data. There is not significant drop in accuracy, it is even slightly better than Dropout in the case of CIFAR-10. This support the idea that dropout can be seen as data augmentation. Furthermore, it suggests that local dynamics that could avoid co-adaptation of neurons may not play a significant role in the training of the networks.

In figure 4 we plotted the intermediate layer representation of a 1000-2000-3000 hidden unit network trained with dropout corruption on CIFAR-10. We computed the intermediate layer representation of four different inputs⁴, i.e., a training sample x_1 , corrupted version of it termed \tilde{x}_1 , x_2 , another sample from the same class as that of x_1 and a sample from different class x_3 . It can be observed from the figure that the representations of corrupted input \tilde{x}_1 and the one from the same class x_2 differ in a similar way from the representation of x_1 while being a lot different from that of x_3 . This suggests that **certain level corruption in hidden layers is similar to sampling a new input from the same class.**

6 Conclusion

In this work we suggest interpreting multiplicative noise such as dropout as data augmentation without domain knowledge. We gave empirical evidence in favor of this view. Further experiments will be required to better understand how this data augmentation procedure evolves during training, how it scales to large architectures, how it applies to data other than images, etc.

We also discussed the effect of increased sparsity levels found in networks trained with Dropout and showed how this is related to data augmentation via input space tiling. Based on richer data augmentation idea, we tested a random dropout noise scheme which improved Dropout results without adding a significant computational cost.

⁴All the inputs are classified correctly.

Acknowledgements

This work was supported by the German Federal Ministry of Education and Research (BMBF) in the project 01GQ0841 (BFNT Frankfurt), an NSERC Discovery grant, a Google faculty research award and Ubisoft. We are grateful towards the developers of Theano [16, 17].

References

- [1] Kurt Wiesenfeld, Frank Moss, et al. Stochastic resonance and the benefits of noise: from ice ages to crayfish and squids. *Nature*, 373(6509):33–36, 1995.
- [2] Yosef Yarom and Jorn Hounsgaard. Voltage fluctuations in neurons: signal or noise? *Physiological reviews*, 91(3):917–929, 2011.
- [3] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008.
- [4] Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012.
- [5] Li Wan, Matthew Zeiler, Sixin Zhang, Yann L Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1058–1066, 2013.
- [6] Pierre Baldi and Peter J Sadowski. Understanding dropout. In *Advances in Neural Information Processing Systems*, pages 2814–2822, 2013.
- [7] Kishore Konda, Roland Memisevic, and David Krueger. Zero-bias autoencoders and the benefits of co-adapting features. In *International Conference on Learning Representations (ICLR)*, 2015.
- [8] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [9] Patrice Y Simard, Dave Steinkraus, and John C Platt. Best practices for convolutional neural networks applied to visual document analysis. In *2013 12th International Conference on Document Analysis and Recognition*, volume 2, pages 958–958. IEEE Computer Society, 2003.
- [10] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [11] Dan Ciresan, Ueli Meier, and Jürgen Schmidhuber. Multi-column deep neural networks for image classification. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3642–3649. IEEE, 2012.
- [12] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.
- [13] Krzysztof J Geras and Charles Sutton. Scheduled denoising autoencoders. *arXiv preprint arXiv:1406.3269*, 2014.
- [14] B Chandra and Rajesh Kumar Sharma. Adaptive noise schedule for denoising autoencoder. In *Neural Information Processing*, pages 535–542. Springer, 2014.
- [15] Steven Rennie, Vaibhava Goel, and Samuel Thomas. Annealed dropout training of deep networks. In *Spoken Language Technology (SLT), IEEE Workshop on. IEEE*, 2014.
- [16] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. Theano: a cpu and gpu math expression compiler. In *Proceedings of the Python for scientific computing conference (SciPy)*, volume 4, page 3. Austin, TX, 2010.
- [17] Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian Goodfellow, Arnaud Bergeron, Nicolas Bouchard, David Warde-Farley, and Yoshua Bengio. Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*, 2012.

7 Appendix

7.1 Proof of x^* unlikeliness

We can show, with a proof by contradiction, that it's unlikely to find a single $x^* = x^{(1)*} = x^{(2)*} = \dots = x^{(n)*}$ that minimizes well L .

By the associative property of function composition, we can rewrite equation 10

$$f^{(i)}(x^*) = (a^{(i)} \circ h^{(i)}) (f^{(i-1)}(x^*)) \quad (12)$$

Suppose there exist an x^* such that

$$(a^{(i)} \circ h^{(i)}) (f^{(i-1)}(x^*)) = (\tilde{a}^{(i)} \circ h^{(i)}) (\tilde{f}^{(i-1)}(x)) \quad (13)$$

$$(a^{(i-1)} \circ h^{(i-1)}) (f^{(i-2)}(x^*)) = (\tilde{a}^{(i-1)} \circ h^{(i-1)}) (\tilde{f}^{(i-2)}(x)) \quad (14)$$

Based on 12 and 14, we have that $f^{(i-1)}(x^*) = \tilde{f}^{(i-1)}(x)$. The proof is concluded by replacing the latter in 13 and then expanding the composed functions.

$$\begin{aligned} (a^{(i)} \circ h^{(i)}) (f^{(i-1)}(x^*)) &= (\tilde{a}^{(i)} \circ h^{(i)}) (f^{(i-1)}(x^*)) \\ \text{rect} \left(h^{(i)} \left(f^{(i-1)}(x^*) \right) \right) &= \vec{m}^{(i)} \odot \text{rect} \left(h^{(i)} \left(f^{(i-1)}(x^*) \right) \right) \end{aligned} \quad (15)$$

Equation 15 can only be true if $\vec{m}^{(i)}$ does not apply any modification to $\text{rect} \left(h^{(i)} \left(f^{(i-1)}(x^*) \right) \right)$, that means $\vec{m}_j^{(i)} = 1$ when $\text{rect}_j \left(h^{(i)} \left(f^{(i-1)}(x^*) \right) \right) > 0$. It happens with a probability $p_{(i)}^{d_{(i)}s_{(i)}}$ where $p_{(i)}$ is the Bernoulli success probability, $d_{(i)}$ is the number of hidden units and $s_{(i)}$ is the mean sparsity level, i.e. mean percentage of active hidden units, of the i -th hidden layer. This probability is very low for standard hyper-parameters values. For instance, with $p_{(i)} = 0.5$, $d_{(i)} = 1000$ and $s_{(i)} = 0.15$, the probability is as low as 10^{-47} .

7.2 Sparsity in dropout networks

Section 5(Figure 3) presented sparsity levels corresponding to Random-dropout. In this section we present additional plots corresponding to the dropout validation experiments in Figure 5. It can be observed from the figure that the proportionality between the input noise level and the corresponding hidden layer sparsity also holds in dropout experiments as stated in Section 5. In case of Figure 5(a) the sparsity initially increases with the input noise level but then starts to fall. This may be due to the dense PCA representation of the input data in the CIFAR-10 experiments where higher noise levels probably destroy the input manifold structure.

7.3 Proportional Gaussian Noise (PGN)

It is shown by [5] that a noise scheme involving multiplication with a zero-mask can be approximated as a Gaussian noise distribution on the hidden units:

$$h_n = \vec{w}_n^T \vec{x} + (\vec{w}_n^2)^T \vec{x}^2 \mathcal{N}(0, p(p-1)) + b_n \quad (16)$$

where \vec{w}_n^2 and \vec{x}^2 are element-wise squares. From equation 16 it is clear that the noise variance will span with $(\vec{w}_n^2)^T \vec{x}^2$. Interestingly, the variance is proportional to the intensity of the response. We used equation 16 in a set of experiments rather than a binary mask, but replaced the expression $p(p-1)$ with hyperparameter σ^2 . We call this approach "Proportional Gaussian Noise" (PGN). The hidden response when using PGN can be written as

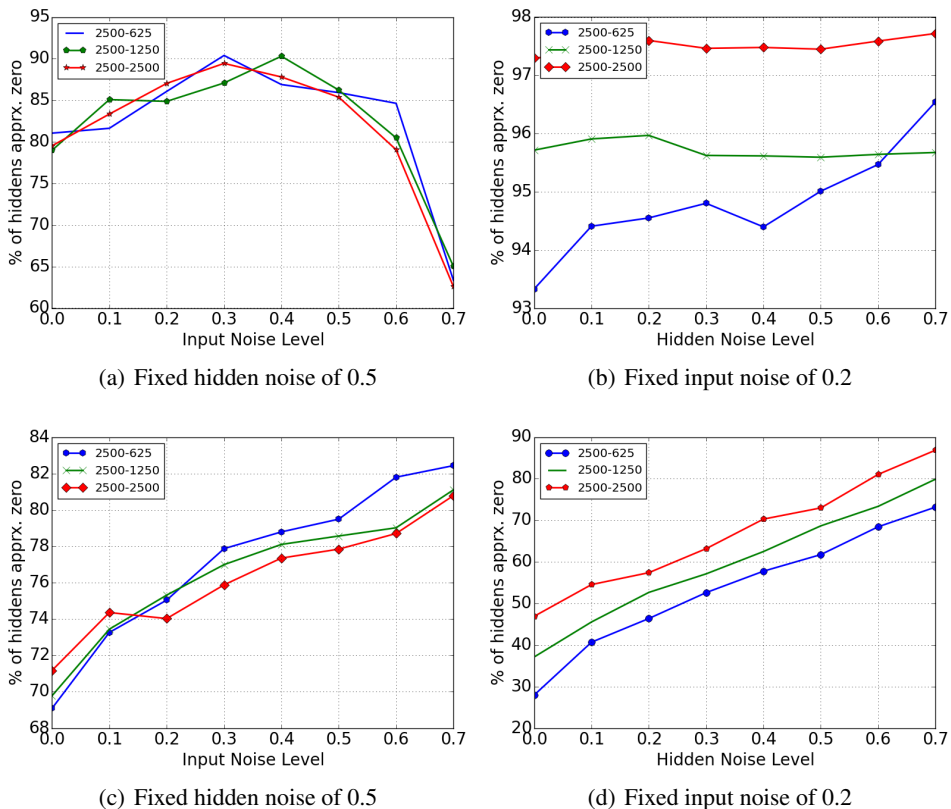


Figure 5: Sparsity in different layers of MLPs trained using noise based regularization schemes. **(a)** and **(b)** correspond to sparsity in layer-1 and layer-2 respectively of networks trained on CIFAR-10 using dropout. Similarly **(d)** and **(e)** are on MNIST.

$$h_n = \vec{w}_n^T \vec{x} + (\vec{w}_n^2)^T \vec{x}^2 \mathcal{N}(0, \sigma^2) + b_n \quad (17)$$

We evaluated the noise scheme PGN where instead of using zero-mask input corruption like in Dropout experiments we add zero mean Gaussian noise with fixed standard deviation to the input. Since noise is added to the input which produces the desired response variation in the first hidden layer units, we do not add any noise to the first layer responses for these experiments. For all the other layers including output layer noise scheme PGN is used. We use same approach as described in Section 5 to train and report PGN network performance. We found that the PGN noise scheme also gives results similar to dropout. A network with 2500-1250 hidden units using PGN noise scheme with input Gaussian noise of 0.2 variance achieved 0.92% error on MNIST dataset.