

# Stochastic Receptive Fields in Deep Convolutional Networks

Audrey G. Chung  
Mohammad Javad Shafiee  
Alexander Wong

University of Waterloo, ON, Canada  
University of Waterloo, ON, Canada  
University of Waterloo, ON, Canada

## Abstract

Deep convolutional neural networks (ConvNets) have rapidly grown in popularity due to their powerful capabilities in representing and modelling the high-level abstraction of complex data. However, ConvNets require an abundance of data to adequately train network parameters. To tackle this problem, we introduce the concept of stochastic receptive fields, where the receptive fields are stochastic realizations of a random field that obey a learned distribution. We study the efficacy of incorporating layers of stochastic receptive fields to a ConvNet to boost performance without the need for additional training data. Preliminary results showing an improvement in accuracy ( $\sim 2\%$  drop in test error) was achieved by adding a layer of stochastic receptive fields to a ConvNet compared to adding a layer of fully-trained receptive fields, when training with a small training set consisting of 20% of the STL-10 dataset.

## 1 Introduction

Deep neural networks (DNNs) are a branch of machine learning that has rapidly grown in popularity as a result of its significant capabilities in representing and modelling high-level abstractions in extremely complex data. DNNs have been shown to perform well in a variety of complex tasks including speech recognition [1], object recognition [2, 3], and natural language processing [4].

A particularly popular form of DNNs is deep convolutional networks (ConvNets). In many of these applications, there is often insufficient training data to adequately train a conventional ConvNet. Data insufficiency is particularly concerning for ConvNets, as they characteristically have a large number of network parameters that require an abundance of labeled training data to learn. Previous works have attempted to address this issue through unsupervised learning of convolutional layers using unlabelled data [5]. While this reduced the amount of labelled data needed, it still requires an abundance of unlabelled data, which may not be available in many situations. This motivates the exploration of alternative strategies for improving the performance of ConvNets in the face of data insufficiency. In this study, we introduce the concept of stochastic receptive fields, where the receptive fields are stochastic realizations of a random field that obey a learned distribution, and study their efficacy for boosting the performance of ConvNets in the situation of small training data sizes.

## 2 Methodology

In traditional ConvNets, performance typically increases with the number of network parameters; however, the amount of training data needed to train the network parameters also increases. The idea behind stochastic receptive fields is to incorporate additional layers of receptive fields that are stochastic realizations of a random field that obey a learned distribution into conventional ConvNets. Since the parameters associated with the distribution used to form stochastic receptive fields in a given layer is significantly fewer than the parameters of fully-trained receptive fields in a given layer, one can potentially improve performance with a marginal increase in parameters, thus mitigating the need for a large increase in training data. We will refer to ConvNets with layers of stochastic receptive fields as RandomNets for the remainder of this paper.

To study the efficacy of RandomNets, we construct a 3-layer RandomNet with 32, 64, and 64  $5 \times 5$  receptive fields in the first, second, and third layers, respectively. The first layer is comprised of stochastic receptive fields formed based on distribution learned parametrically from a set of Gabor-like receptive fields. The second and third layers consist of fully-trained receptive fields. Lastly, there is a fully-connected hidden layer of 64 neurons. This RandomNet is compared to a baseline ConvNet with all three convolutional layers having fully-trained receptive fields.

To assess the performance of RandomNet, we used the STL-10 dataset [6], which contains 5000 labelled training images and 8000 labelled images, categorized into 10 classes. The parameters

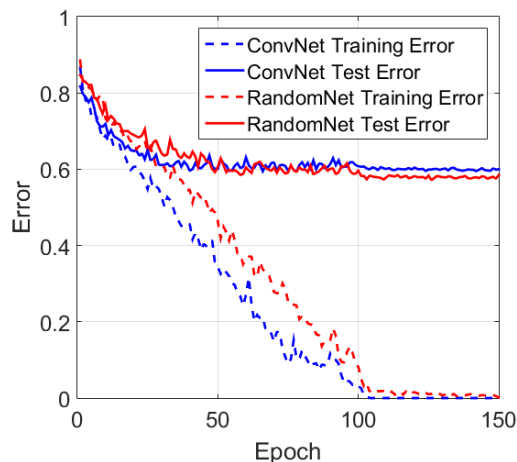


Fig. 1: Training and testing error for STL-10 dataset for the baseline ConvNet (blue) and the proposed RandomNet (red), when trained with 20% of STL-10 training data. The RandomNet has improved performance ( $\sim 2\%$ ) relative to the ConvNet.

of both RandomNet and ConvNet were trained using 20% of the STL-10 dataset; Fig. 1 shows the training and testing errors for the baseline ConvNet and the proposed RandomNet.

## 3 Discussion

It can be observed from Fig. 1 that the replacement of a layer of fully-trained receptive fields with stochastic receptive fields in a ConvNet can indeed improve performance in a situation characterized by small training size. More specifically, the tested RandomNet demonstrated a decrease in test error over the baseline ConvNet by  $\sim 2\%$ , even though the baseline ConvNet has an additional layer of fully-trained receptive fields and correspondingly more network parameters compared to the RandomNet. These preliminary results illustrate the potential of adding layers of stochastic receptive fields for boosting the performance of ConvNets without necessitating additional training data. Further investigation is required to better understand the benefits and limitations of stochastic receptive fields and evaluate its performance in various machine learning applications.

## Acknowledgments

This research has been supported by the Canada Research Chairs Program, Natural Sciences and Engineering Research Council of Canada (NSERC), and the Ministry of Research and Innovation of Ontario. The authors also thank Nvidia for the GPU hardware used in this study through the Nvidia Hardware Grant Program.

## References

- [1] Hannun, A. et al. Deep Speech: Scaling Up End-to-End Speech Recognition. *arXiv* 1-12 (2014).
- [2] Krizhevsky, A., Sutskever, I. and Hinton, G. Imagenet classification with deep convolutional neural networks. *NIPS* (2012).
- [3] Shafiee, M. J., Siva, P. and Wong, A. StochasticNet: Forming Deep Neural Networks via Stochastic Connectivity. *arXiv* (2015).
- [4] Bengio, Y. et al. A neural probabilistic language model. *JMLR* 3 1137-1155 (2003).
- [5] Hadsell, R. et al. A. Learning long-range vision for autonomous off-road driving. *Journal of Field Robotics* (2009).
- [6] Coates, A., Ng, A. and Lee, H. An Analysis of Single Layer Networks in Unsupervised Features Learning. *AISTATS* (2011).