

DEEP CONVOLUTIONAL NEURAL NETWORK DESIGN PATTERNS

Leslie N. Smith

U.S. Naval Research Laboratory, Code 5514
4555 Overlook Ave., SW., Washington, D.C. 20375
leslie.smith@nrl.navy.mil

Nicholay Topin

University of Maryland, Baltimore County
1000 Hilltop Circle, Baltimore, MD 21250
ntopin1@umbc.edu

ABSTRACT

Recent research in the deep learning field has produced a plethora of new architectures. At the same time, a growing number of groups are applying deep learning to new applications. Some of these groups are likely to be composed of inexperienced deep learning practitioners who are baffled by the dizzying array of architecture choices and therefore opt to use an older architecture (i.e., Alexnet). Here we attempt to bridge this gap by mining the collective knowledge contained in recent deep learning research to discover underlying principles for designing neural network architectures. In addition, we describe several architectural innovations, including Fractal of FractalNet network, Stagewise Boosting Networks, and Taylor Series Networks (our Caffe code and prototxt files is available at <https://github.com/iPhysicist/CNNDesignPatterns>). We hope others are inspired to build on our preliminary work.

1 INTRODUCTION

Many recent articles discuss new architectures for neural networking, especially regarding Residual Networks (He et al. (2015; 2016); Larsson et al. (2016); Zhang et al. (2016); Huang et al. (2016b)). Although the literature covers a wide range of network architectures, we take a high-level view of the architectures as the basis for discovering universal principles of the design of network architecture. We discuss 14 original design patterns that could benefit inexperienced practitioners who seek to incorporate deep learning in various new applications. This paper addresses the current lack of guidance on design, a deficiency that may prompt the novice to rely on standard architecture, e.g., Alexnet, regardless of the architecture’s suitability to the application at hand.

This abundance of research is also an opportunity to determine which elements provide benefits in what specific contexts. We ask: Do universal principles of deep network design exist? Can these principles be mined from the collective knowledge on deep learning? Which architectural choices work best in any given context? Which architectures or parts of architectures seem elegant?

Design patterns were first described by Christopher Alexander (Alexander (1979)) in regards to the architectures of buildings and towns. Alexander wrote of a timeless quality in architecture that “lives” and this quality is enabled by building based on universal principles. The basis of design patterns is that they resolve a conflict of forces in a given context and lead to an equilibrium analogous to the ecological balance in nature. Design patterns are both highly specific, making them clear to follow, and flexible so they can be adapted to different environments and situations. Inspired by Alexander’s work, the “gang of four” (Gamma et al. (1995)) applied the concept of design patterns to the architecture of object-oriented software. This classic computer science book describes 23 patterns that resolve issues prevalent in software design, such as “requirements always change”. We were inspired by these previous works on architectures to articulate possible design patterns for convolutional neural network (CNN) architectures.

Design patterns provide universal guiding principles, and here we take the first steps to defining network design patterns. Overall, it is an enormous task to define design principles for all neural networks and all applications, so we limit this paper to CNNs and their canonical application of image classification. However, we recognize that architectures must depend upon the application by defining our first design pattern; *Design Pattern 1: Architectural Structure follows the Application*

(we leave the details of this pattern to future work). In addition, these principles allowed us to discover some gaps in the existing research and to articulate novel networks (i.e, Fractal of FractalNets, Stagemwise Boosting and Taylor Series networks) and units (i.e., freeze-drop-path). We hope the rules of thumb articulated here are valuable for both the experienced and novice practitioners and that our preliminary work serves as a stepping stone for others to discover and share additional deep learning design patterns.

2 RELATED WORK

To the best of our knowledge, there has been little written recently to provide guidance and understanding on appropriate architectural choices¹. The book *"Neural Networks: Tricks of the Trade"* (Orr & Müller, 2003) contains recommendations for network models but without reference to the vast amount of research in the past few years. Perhaps the closest to our work is Szegedy et al. (2015b) where those authors describe a few design principles based on their experiences.

Much research has studied neural network architectures, but we are unaware of a recent survey of the field. Unfortunately, we cannot do justice to this entire body of work, so we focus on recent innovations in convolutional neural networks architectures and, in particular, on Residual Networks (He et al., 2015) and its recent family of variants. We start with Network In Networks (Lin et al., 2013), which describes a hierarchical network with a small network design repeatedly embedded in the overall architecture. Szegedy et al. (2015a) incorporated this idea into their Inception architecture. Later, these authors proposed modifications to the original Inception design (Szegedy et al., 2015b). A similar concept was contained in the multi-scale convolution architecture (Liao & Carneiro, 2015). In the meantime, Batch Normalization (Ioffe & Szegedy, 2015) was presented as a unit within the network that makes training faster and easier.

Before the introduction of Residual Networks, a few papers suggested skip connections. Skip connections were proposed by Raiko et al. (2012). Highway Networks (Srivastava et al., 2015) use a gating mechanism to decide whether to combine the input with the layer's output and showed how these networks allowed the training of very deep networks. The DropIn technique (Smith et al., 2015; 2016) also trains very deep networks by allowing a layer's input to skip the layer. The concept of stochastic depth via a drop-path method was introduced by Huang et al. (2016b).

Residual Networks were introduced by He et al. (2015), where the authors describe their network that won the 2015 ImageNet Challenge. They were able to extend the depth of a network from tens to hundreds of layers and in doing so, improve the network's performance. The authors followed up with another paper (He et al., 2016) where they investigate why identity mappings help and report results for a network with more than a thousand layers. The research community took notice of this architecture and many modifications to the original design were soon proposed.

The Inception-v4 paper (Szegedy et al., 2016) describes the impact of residual connections on their Inception architecture and compared these results with the results from an updated Inception design. The Resnet in Resnet paper (Targ et al., 2016) suggests a duel stream architecture. Veit et al. (2016) provided an understanding of Residual Networks as an ensemble of relatively shallow networks. These authors illustrated how these residual connections allow the input to follow an exponential number of paths through the architecture. At the same time, the FractalNet paper (Larsson et al., 2016) demonstrated training deep networks with a symmetrically repeating architectural pattern. As described later, we found the symmetry introduced in their paper intriguing. In a similar vein, Convolutional Neural Fabrics (Saxena & Verbeek, 2016) introduces a three dimensional network, where the usual depth through the network is the first dimension.

Wide Residual Networks (Zagoruyko & Komodakis, 2016) demonstrate that simultaneously increasing both depth and width leads to improved performance. In Swapout (Singh et al., 2016), each layer can be dropped, skipped, used normally, or combined with a residual. Deeply Fused Nets (Wang et al., 2016) proposes networks with multiple paths. In the Weighted Residual Networks paper (Shen & Zeng, 2016), the authors recommend a weighting factor for the output from the convolutional layers, which gradually introduces the trainable layers. Convolutional Residual Memory Networks (Moniz & Pal, 2016) proposes an architecture that combines a convolutional Residual Network with

¹After submission we became aware of an online book being written on deep learning design patterns at <http://www.deeplearningpatterns.com>

an LSTM memory mechanism. For Residual of Residual Networks (Zhang et al., 2016), the authors propose adding a hierarchy of skip connections where the input can skip a layer, a module, or any number of modules. DenseNets (Huang et al., 2016a) introduces a network where each module is densely connected; that is, the output from a layer is input to all of the other layers in the module. In the Multi-Residual paper (Abdi & Nahavandi, 2016), the authors propose expanding a residual block width-wise to contain multiple convolutional paths. Our Appendix A describes the close relationship between many of these Residual Network variants.

3 DESIGN PATTERNS

We reviewed the literature specifically to extract commonalities and reduce their designs down to fundamental elements that might be considered design patterns. It seemed clear to us that in reviewing the literature some design choices are elegant while others are less so. In particular, the patterns described in this paper are the following:

1. *Architectural Structure follows the Application*
2. *Proliferate Paths*
3. *Strive for Simplicity*
4. *Increase Symmetry*
5. *Pyramid Shape*
6. *Over-train*
7. *Cover the Problem Space*
8. *Incremental Feature Construction*
9. *Normalize Layer Inputs*
10. *Input Transition*
11. *Available Resources Guide Layer Widths*
12. *Summation Joining*
13. *Down-sampling Transition*
14. *Maxout for Competition*

3.1 HIGH LEVEL ARCHITECTURE DESIGN

Several researchers have pointed out that the winners of the ImageNet Challenge (Russakovsky et al., 2015) have successively used deeper networks (as seen in, Krizhevsky et al. (2012), Szegedy et al. (2015a), Simonyan & Zisserman (2014), He et al. (2015)). It is also apparent to us from the ImageNet Challenge that multiplying the number of paths through the network is a recent trend that is illustrated in the progression from Alexnet to Inception to ResNets. For example, Veit et al. (2016) show that ResNets can be considered to be an exponential ensemble of networks with different lengths. *Design Pattern 2: Proliferate Paths* is based on the idea that ResNets can be an exponential ensemble of networks with different lengths. One proliferates paths by including a multiplicity of branches in the architecture. Recent examples include FractalNet (Larsson et al. 2016), Xception (Chollet 2016), and Decision Forest Convolutional Networks (Ioannou et al. 2016).

Scientists have embraced simplicity/parsimony for centuries. Simplicity was exemplified in the paper "Striving for Simplicity" (Springenberg et al. 2014) by achieving state-of-the-art results with fewer types of units. *Design Pattern 3: Strive for Simplicity* suggests using fewer types of units and keeping the network as simple as possible. We also noted a special degree of elegance in the FractalNet (Larsson et al. 2016) design, which we attributed to the symmetry of its structure. *Design Pattern 4: Increase Symmetry* is derived from the fact that architectural symmetry is typically considered a sign of beauty and quality. In addition to its symmetry, FractalNets also adheres to the *Proliferate Paths* design pattern so we used it as the baseline of our experiments in Section 4.

An essential element of design patterns is the examination of trade-offs in an effort to understand the relevant forces. One fundamental trade-off is the maximization of representational power versus

elimination of redundant and non-discriminating information. It is universal in all convolutional neural networks that the activations are downsampled and the number of channels increased from the input to the final layer, which is exemplified in Deep Pyramidal Residual Networks (Han et al. (2016)). *Design Pattern 5: Pyramid Shape* says there should be an overall smooth downsampling combined with an increase in the number of channels throughout the architecture.

Another trade-off in deep learning is training accuracy versus the ability of the network to generalize to non-seen cases. The ability to generalize is an important virtue of deep neural networks. Regularization is commonly used to improve generalization, which includes methods such as dropout (Srivastava et al. 2014a) and drop-path (Huang et al. 2016b). As noted by Srivastava et al. 2014b, dropout improves generalization by injecting noise in the architecture. We believe regularization techniques and prudent noise injection during training improves generalization (Srivastava et al. 2014b, Gulcehre et al. 2016). *Design Pattern 6: Over-train* includes any training method where the network is trained on a harder problem than necessary to improve generalization performance of inference. *Design Pattern 7: Cover the Problem Space* with the training data is another way to improve generalization (e.g., Ratner et al. 2016, Hu et al. 2016, Wong et al. 2016, Johnson-Roberson et al. 2016). Related to regularization methods, cover the problem space includes the use of noise (Rasmus et al. 2015, Krause et al. 2015, Pezeshki et al. 2015), synthetic data, and data augmentation, such as random cropping, flipping, and varying brightness, contrast, and the like.

3.2 DETAILED ARCHITECTURE DESIGN

A common thread throughout many of the more successful architectures is to make each layer’s “job” easier. Use of very deep networks is an example because any single layer only needs to incrementally modify the input, and this partially explains the success of Residual Networks, since in very deep networks, a layer’s output is likely similar to the input; hence adding the input to the layer’s output makes the layer’s job incremental. Also, this concept is part of the motivation behind design pattern 2 but it extends beyond that. *Design Pattern 8: Incremental Feature Construction* recommends using short skip lengths in ResNets. A recent paper (Alain & Bengio (2016)) showed in an experiment that using an identity skip length of 64 in a network of depth 128 led to the first portion of the network not being trained.

Design Pattern 9: Normalize Layer Inputs is another way to make a layer’s job easier. Normalization of layer inputs has been shown to improve training and accuracy but the underlying reasons are not clear (Ioffe & Szegedy 2015, Ba et al. 2016, Salimans & Kingma 2016). The Batch Normalization paper (Ioffe & Szegedy 2015) attributes the benefits to handling internal covariate shift, while the authors of streaming normalization (Liao et al. 2016) express that it might be otherwise. We feel that normalization puts all the layer’s input samples on more equal footing (analogous to a units conversion scaling), which allows back-propagation to train more effectively.

Some research, such as Wide ResNets (Zagoruyko & Komodakis 2016), has shown that increasing the number of channels improves performance but there are additional costs with extra channels. The input data for many of the benchmark datasets have 3 channels (i.e., RGB). *Design Pattern 10: Input Transition* is based on the common occurrence that the output from the first layer of a CNN significantly increases the number of channels from 3. A few examples of this increase in channels/outputs at the first layer for ImageNet are AlexNet (96 channels), Inception (32), VGG (224), and ResNets (64). Intuitively it makes sense to increase the number of channels from 3 in the first layer as it allows the input data to be examined many ways but it is not clear how many outputs are best. Here, the trade-off is that of cost versus accuracy. Costs include the number of parameters in the network, which directly affects the computational and storage costs of training and inference. *Design Pattern 11: Available Resources Guide Layer Widths* is based on balancing costs against an application’s requirements. Choose the number of outputs of the first layer based on memory and computational resources and desired accuracy.

3.2.1 JOINING BRANCHES: CONCATENATION, SUMMATION/MEAN, MAXOUT

When there are multiple branches, three methods have been used to combine the outputs; concatenation, summation (or mean), or Maxout. It seems that different researchers have their favorites and there hasn’t been any motivation for using one in preference to another. In this Section, we propose some rules for deciding how to combine branches.

Summation is one of the most common ways of combining branches. *Design Pattern 12: Summation Joining* is where the joining is performed by summation/mean. Summation is the preferred joining mechanism for Residual Networks because it allows each branch to compute corrective terms (i.e., residuals) rather than the entire approximation. The difference between summation and mean (i.e., fractal-join) is best understood by considering drop-path (Huang et al. 2016b). In a Residual Network where the input skip connection is always present, summation causes the layers to learn the residual (the difference from the input). On the other hand, in networks with several branches, where any branch can be dropped (e.g., FractalNet (Larsson et al. (2016))), using the mean is preferable as it keeps the output smooth as branches are randomly dropped.

Some researchers seem to prefer concatenation (e.g., Szegedy et al. (2015a)). *Design Pattern 13: Down-sampling Transition* recommends using concatenation joining for increasing the number of outputs when pooling. That is, when down-sampling by pooling or using a stride greater than 1, a good way to combine branches is to concatenate the output channels, hence smoothly accomplishing both joining and an increase in the number of channels that typically accompanies down-sampling.

Maxout has been used for competition, as in locally competitive networks (Srivastava et al. 2014b) and competitive multi-scale networks Liao & Carneiro (2015). *Design Pattern 14: Maxout for Competition* is based on Maxout choosing only one of the activations, which is in contrast to summation or mean where the activations are “cooperating”; here, there is a “competition” with only one “winner”. For example, when each branch is composed of different sized kernels, Maxout is useful for incorporating scale invariance in an analogous way to how max pooling enables translation invariance.

4 EXPERIMENTS

4.1 ARCHITECTURAL INNOVATIONS

In elucidating these fundamental design principles, we also discovered a few architectural innovations. In this section we will describe these innovations.

First, we recommended combining summation/mean, concatenation, and maxout joining mechanisms with differing roles within a single architecture, rather than the typical situation where only one is used throughout. Next, *Design Pattern 2: Proliferate Branches* led us to modify the overall sequential pattern of modules in the FractalNet architecture. Instead of lining up the modules for maximum depth, we arranged the modules in a fractal pattern as shown in 1b, which we named a Fractal of FractalNet (FoF) network, where we exchange depth for a greater number of paths.

4.1.1 FREEZE-DROP-PATH AND STAGewise BOOSTING NETWORKS (SBN)

Drop-path was introduced by Huang et al. (2016b), which works by randomly removing branches during an iteration of training, as though that path doesn’t exist in the network. Symmetry considerations led us to an opposite method that we named freeze-path. Instead of removing a branch from the network during training, we freeze the weights, as though the learning rate was set to zero. A similar idea has been proposed for recurrent neural networks (Krueger et al. 2016).

The potential usefulness of combining drop-path and freeze-path, which we named freeze-drop-path, is best explained in the non-stochastic case. Figure 1 shows an example of a fractal of FractalNet architecture. Let’s say we start training only the leftmost branch in Figure 1b and use drop-path on all of the other branches. This branch should train quickly since it has only a relatively few parameters compared to the entire network. Next we freeze the weights in that branch and allow the next branch to the right to be active. If the leftmost branch is providing a good function approximation, the next branch works to produce a “small” corrective term. Since the next branch contains more layers than the previous branch and the corrective term should be easier to approximate than the original function, the network should attain greater accuracy. One can continue this process from left to right to train the entire network. We used freeze-drop-path as the final/bottom join in the FoF architecture in Figure 1b and named this the Stagewise Boosting Networks (SBN) because they are analogous to stagewise boosting (Friedman et al. 2001). The idea of boosting neural networks is not new (Schwenk & Bengio 2000) but this architecture is new. In Appendix B we discuss the implementation we tested.

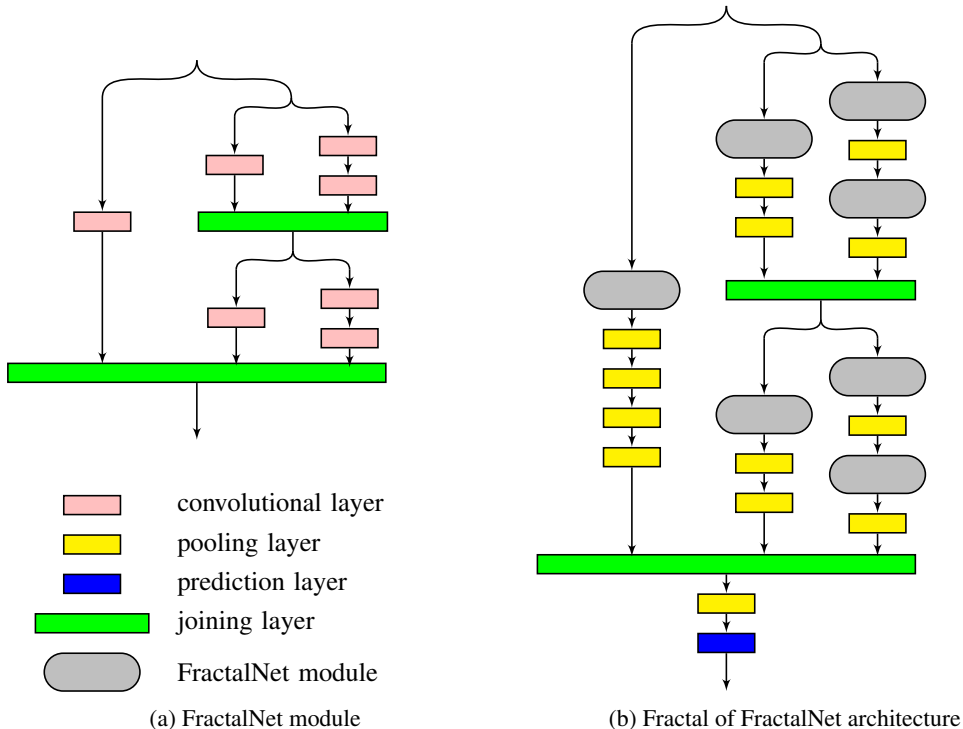


Figure 1: (a) The FractalNet module and (b) the FoF architecture.

4.1.2 TAYLOR SERIES NETWORKS (TSN)

Taylor series expansions are classic and well known as a function approximation method, which is:

$$f(x+h) = f(x) + hf'(x) + h^2 f''(x)/2 + \dots \quad (1)$$

Since neural networks are also function approximators, it is a short leap from FoFs and SBNs to consider the branches of that network as terms in a Taylor series expansion. Hence, the Taylor series implies squaring the second branch before the summation joining unit, analogous to the second order term in the expansion. Similarly, the third branch would be cubed. We call this ‘‘Taylor Series Networks’’ (TSN) and there is precedence for this idea in the literature with polynomial networks (Livni et al. 2014) and multiplication in networks (e.g. Lin et al. 2015). The implementation details of this TSN are discussed in the Appendix.

Table 1: Comparison of test accuracy results at the end of the training, averaged over three runs.

Architecture	CIFAR-10 (%)	CIFAR-100 (%)
FractalNet	93.28 ± 0.22	72.39 ± 0.10
FractalNet + Concat	93.05 ± 0.05	72.57 ± 0.31
FractalNet + Maxout	91.99 ± 0.33	69.98 ± 0.24
FractalNet + Avg pooling	94.25 ± 0.09	72.81 ± 0.52
FoF	92.31 ± 0.25	72.81 ± 0.28
SBN	90.69 ± 0.62	68.70 ± 0.86
TSN	90.7 ± 0.04	68.36 ± 0.10

4.2 RESULTS

The experiments in this section are primarily to empirically validate the architectural innovations described above but not to fully test them. We leave a more complete evaluation to future work.

Table 1 and Figures 2 and 3 compare the final test accuracy results for CIFAR-10 and CIFAR-100 in a number of experiments. An accuracy value in Table 1 is computed as the mean of the last 6 test




Figure 2: Test accuracy of the original FractalNet compared with replacing some of the fractal-joints with Concatenation or Maxout, and when replacing max pooling with average pooling.




Figure 3: Test accuracy of the original FractalNet compared with FoF, SBN, and TSN networks.

accuracies computed over the last 3,000 iterations (out of 100,000) of the training. The results from the original FractalNet (Larsson et al. 2016) are given in the first row of the table and we use this as our baseline. The first four rows of Table 1 and Figure 2 compare the test accuracy of the original FractalNet architectures to architectures with a few modifications advocated by design patterns. The first modification is to use concatenation instead of fractal-joints at all the downsampling locations in the networks. The results for both CIFAR-10 (2a) and CIFAR-100 (2b) indicate that the results are equivalent when concatenation is used instead of fractal-joints at all the downsampling locations in the networks. The second experiment was to change the kernel sizes in the first module from 3x3 such that the left most column used a kernel size of 7x7, the second column 5x5, and the third used 3x3. The fractal-join for module one was replaced with Maxout. The results in Figure 2 are a bit worse, indicating that the more cooperative fractal-join (i.e., mean/summation) with 3x3 kernels has better performance than the competitive Maxout with multiple scales. Figure 2 also illustrates how an experiment replacing max pooling with average pooling throughout the architecture changes the training profile. For CIFAR-10, the training accuracy rises quickly, plateaus, then lags behind the original FractalNet but ends with a better final performance, which implies that average pooling might significantly reduce the length of the training (we plan to examine this in future work). This behavior provides some evidence that “cooperative” average pooling might be preferable to “competitive” max pooling.

Table 1 and Figure 3 compare the test accuracy results for the architectural innovations described in Section 4.1. The FoF architecture ends with a similar final test accuracy as FractalNet but the SBN and TSN architectures (which use freeze-drop-path) lag behind when the learning rate is dropped. However, it is clear from both Figures 3a and 3b that the new architectures train more quickly than FractalNet. The FoF network is best as it trains more quickly than FractalNet and achieves similar accuracy. The use of freeze-drop-path in SBN and TSN is questionable since the final performance lags behind FractalNet, but we are leaving the exploration for more suitable applications of these new architectures for future work.

5 CONCLUSION

In this paper we describe convolutional neural network architectural design patterns that we discovered by studying the plethora of new architectures in recent deep learning papers. We hope these design patterns will be useful to both experienced practitioners looking to push the state-of-the-art and novice practitioners looking to apply deep learning to new applications. There exists a large expanse of potential follow up work (some of which we have indicated here as future work). Our effort here is primarily focused on Residual Networks for classification but we hope this preliminary work will inspire others to follow up with new architectural design patterns for Recurrent Neural Networks, Deep Reinforcement Learning architectures, and beyond.

ACKNOWLEDGMENTS

The authors want to thank the numerous researchers upon whose work these design patterns are based and especially Larsson et al. 2016 for making their code publicly available. This work was supported by the US Naval Research Laboratory base program.

REFERENCES

- Masoud Abdi and Saeid Nahavandi. Multi-residual networks. *arXiv preprint arXiv:1609.05672*, 2016.
- Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. *arXiv preprint arXiv:1610.01644*, 2016.
- Christopher Alexander. *The timeless way of building*, volume 1. New York: Oxford University Press, 1979.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- François Chollet. Xception: Deep learning with depthwise separable convolution. *arXiv preprint arXiv:1610.02357*, 2016.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.
- Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design patterns: elements of reusable object-oriented software*. Pearson Education India, 1995.
- Caglar Gulcehre, Marcin Moczulski, Misha Denil, and Yoshua Bengio. Noisy activation functions. *arXiv preprint arXiv:1603.00391*, 2016.
- Dongyoon Han, Jiwhan Kim, and Junmo Kim. Deep pyramidal residual networks. *arXiv preprint arXiv:1610.02915*, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. *arXiv preprint arXiv:1603.05027*, 2016.

- Guosheng Hu, Xiaojiang Peng, Yongxin Yang, Timothy Hospedales, and Jakob Verbeek. Frankenstein: Learning deep face representations using small data. *arXiv preprint arXiv:1603.06470*, 2016.
- Gao Huang, Zhuang Liu, and Kilian Q Weinberger. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2016a.
- Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Weinberger. Deep networks with stochastic depth. *arXiv preprint arXiv:1603.09382*, 2016b.
- Yani Ioannou, Duncan Robertson, Darko Zikic, Peter Kotschieder, Jamie Shotton, Matthew Brown, and Antonio Criminisi. Decision forests, convolutional networks and the models in-between. *arXiv preprint arXiv:1603.01250*, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, pp. 675–678. ACM, 2014.
- Matthew Johnson-Roberson, Charles Barto, Rounak Mehta, Sharath Nittur Sridhar, and Ram Vasudevan. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? *arXiv preprint arXiv:1610.01983*, 2016.
- Jonathan Krause, Benjamin Sapp, Andrew Howard, Howard Zhou, Alexander Toshev, Tom Duerig, James Philbin, and Li Fei-Fei. The unreasonable effectiveness of noisy data for fine-grained recognition. *arXiv preprint arXiv:1511.06789*, 2015.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- David Krueger, Tegan Maharaj, János Kramár, Mohammad Pezeshki, Nicolas Ballas, Nan Rosemary Ke, Anirudh Goyal, Yoshua Bengio, Hugo Larochelle, Aaron Courville, et al. Zoneout: Regularizing rnns by randomly preserving hidden activations. *arXiv preprint arXiv:1606.01305*, 2016.
- Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Fractalnet: Ultra-deep neural networks without residuals. *arXiv preprint arXiv:1605.07648*, 2016.
- Qianli Liao, Kenji Kawaguchi, and Tomaso Poggio. Streaming normalization: Towards simpler and more biologically-plausible normalizations for online and recurrent learning. *arXiv preprint arXiv:1610.06160*, 2016.
- Zhibin Liao and Gustavo Carneiro. Competitive multi-scale convolution. *arXiv preprint arXiv:1511.05635*, 2015.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
- Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1449–1457, 2015.
- Roi Livni, Shai Shalev-Shwartz, and Ohad Shamir. On the computational efficiency of training neural networks. In *Advances in Neural Information Processing Systems*, pp. 855–863, 2014.
- Joel Moniz and Christopher Pal. Convolutional residual memory networks. *arXiv preprint arXiv:1606.05262*, 2016.

- Genevieve B Orr and Klaus-Robert Müller. *Neural networks: tricks of the trade*. Springer, 2003.
- Mohammad Pezeshki, Linxi Fan, Philemon Brakel, Aaron Courville, and Yoshua Bengio. Deconstructing the ladder network architecture. *arXiv preprint arXiv:1511.06430*, 2015.
- Tapani Raiko, Harri Valpola, and Yann LeCun. Deep learning made easier by linear transformations in perceptrons. In *AISTATS*, volume 22, pp. 924–932, 2012.
- Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *Advances in Neural Information Processing Systems*, pp. 3546–3554, 2015.
- Alexander Ratner, Christopher De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. Data programming: Creating large training sets, quickly. *arXiv preprint arXiv:1605.07723*, 2016.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- Tim Salimans and Diederik P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *arXiv preprint arXiv:1602.07868*, 2016.
- Shreyas Saxena and Jakob Verbeek. Convolutional neural fabrics. *arXiv preprint arXiv:1606.02492*, 2016.
- Holger Schwenk and Yoshua Bengio. Boosting neural networks. *Neural Computation*, 12(8):1869–1887, 2000.
- Falong Shen and Gang Zeng. Weighted residuals for very deep networks. *arXiv preprint arXiv:1605.08831*, 2016.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Saurabh Singh, Derek Hoiem, and David Forsyth. Swapout: Learning an ensemble of deep architectures. *arXiv preprint arXiv:1605.06465*, 2016.
- Leslie N Smith, Emily M Hand, and Timothy Doster. Gradual dropin of layers to train very deep neural networks. *arXiv preprint arXiv:1511.06951*, 2015.
- Leslie N Smith, Emily M Hand, and Timothy Doster. Gradual dropin of layers to train very deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2016.
- Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014a.
- Rupesh K Srivastava, Klaus Greff, and Jürgen Schmidhuber. Training very deep networks. In *Advances in neural information processing systems*, pp. 2377–2385, 2015.
- Rupesh Kumar Srivastava, Jonathan Masci, Faustino Gomez, and Jürgen Schmidhuber. Understanding locally competitive networks. *arXiv preprint arXiv:1410.1165*, 2014b.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2015a.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Re-thinking the inception architecture for computer vision. *arXiv preprint arXiv:1512.00567*, 2015b.

Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *arXiv preprint arXiv:1602.07261*, 2016.

Sasha Targ, Diogo Almeida, and Kevin Lyman. Resnet in resnet: Generalizing residual architectures. *arXiv preprint arXiv:1603.08029*, 2016.

Andreas Veit, Michael Wilber, and Serge Belongie. Residual networks are exponential ensembles of relatively shallow networks. *arXiv preprint arXiv:1605.06431*, 2016.

Jingdong Wang, Zhen Wei, Ting Zhang, and Wenjun Zeng. Deeply-fused nets. *arXiv preprint arXiv:1605.07716*, 2016.

Sebastien C Wong, Adam Gatt, Victor Stamatescu, and Mark D McDonnell. Understanding data augmentation for classification: when to warp? *arXiv preprint arXiv:1609.08764*, 2016.

Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

Ke Zhang, Miao Sun, Tony X Han, Xingfang Yuan, Liru Guo, and Tao Liu. Residual networks of residual networks: Multilevel residual networks. *arXiv preprint arXiv:1608.02908*, 2016.

A RELATIONSHIPS BETWEEN RESIDUAL ARCHITECTURES

The architectures mentioned in Section 2 commonly combine outputs from two or more layers using concatenation along the depth axis, element-wise summation, and element-wise average. We show here that the latter two are special cases of the former with weight-sharing enforced. Likewise, we show that skip connections can be considered as introducing additional layers into a network that share parameters with existing layers. In this way, any of the Residual Network variants can be reformulated into a standard form where many of the variants are equivalent.

A filter has three dimensions: two spatial dimensions, along which convolution occurs, and a third dimension, depth. Each input channel corresponds to a different depth for each filter of a layer. As a result, a filter can be considered to consist of “slices,” each of which is convolved over one input channel. The results of these convolutions are then added together, along with a bias, to produce a single output channel. The output channels of multiple filters are concatenated to produce the output of a single layer. When the outputs of several layers are concatenated, the behavior is similar to that of a single layer. However, instead of each filter having the same spatial dimensions, stride, and padding, each filter may have a different structure. As far as the function within a network, though, the two cases are the same. In fact, a standard layer, one where all filters have the same shape, can be considered a special case of concatenating outputs of multiple layer types.

If summation is used instead of concatenation, the network can be considered to perform concatenation but enforce weight-sharing in the following layer. The results of first summing several channels element-wise and then convolving a filter slice over the output yields the same result as convolving the slice over the channels and then performing an element-wise summation afterwards. Therefore, enforcing weight-sharing such that the filter slices applied to the n th channel of all inputs share weight results in behavior identical to summation, but in a form similar to concatenation, which highlights the relationship between the two. When Batch Normalization (BN) (Ioffe & Szegedy 2015) is used, as is the current standard practice, performing an average is essentially identical to performing a summation, since BN scales the output. Therefore, scaling the input by a constant (i.e., averaging instead of a summation) is rendered irrelevant. The details of architecture-specific manipulations of summations and averages is described further in Section 3.2.1.

Due to the ability to express depth-wise concatenation, element-wise sum, and element-wise mean as variants of each other, architectural features of recent works can be combined within a single network, regardless of choice of combining operation. However, this is not to say that concatenation has the most expressivity and is therefore strictly better than the others. Summation allows networks to divide up the network’s task. Also, there is a trade-off between the number of parameters and the expressivity of a layer; summation uses weight-sharing to significantly reduce the number of parameters within a layer at the expense of some amount of expressivity.

Different architectures can further be expressed in a similar fashion through changes in the connections themselves. A densely connected series of layers can be “pruned” to resemble any desired architecture with skip connections through zeroing specific filter slices. This operation removes the dependency of the output on a specific input channel; if this is done for all channels from a given layer, the connection between the two layers is severed. Likewise, densely connected layers can be turned into linearly connected layers while preserving the layer dependencies; a skip connection can be passed through the intermediate layers. A new filter can be introduced for each input channel passing through, where the filter performs the identity operation for the given input channel. All existing filters in the intermediate layers can have zeroed slices for this input so as to not introduce new dependencies. In this way, arbitrarily connected layers can be turned into a standard form.

We certainly do not recommend this representation for actual experimentation as it introduces fixed parameters. We merely describe it to illustrate the relationship between different architectures. This representation illustrates how skip connections effectively enforce specific weights in intermediate layers. Though this restriction reduces expressivity, the number of stored weights is reduced, the number of computations performed is decreased, and the network might be more easily trainable.

B IMPLEMENTATION DETAILS

Our implementations are in Caffe (Jia et al. 2014; downloaded October 9, 2016) using CUDA 8.0. These experiments were run on a 64 node cluster with 8 Nvidia Titan Black GPUs, 128 GB memory, and dual Intel Xenon E5-2620 v2 CPUs per node. We used the CIFAR-10 and CIFAR-100 datasets (Krizhevsky & Hinton 2009 for our classification tests. These datasets consist of 60,000 32x32 colour images (50,000 for training and 10,000 for testing) in 10 or 100 classes, respectively. Our Caffe code and prototxt files are publicly available at <https://github.com/iPhysicist/CNNDesignPatterns>.

B.1 ARCHITECTURES

We started with the FractalNet implementation² as our baseline and it is described in Larsson et al. 2016. We used the three column module as shown in Figure 1a. In some of our experiments, we replaced the fractal-join with concatenation at the downsampling locations. In other experiments, we modified the kernel sizes in module one and combined the branches with Maxout. A FractalNet module is shown in Figure 1a and the architecture consists of five sequential modules.

Our fractal of FractalNet (FoF) architecture uses the same module but has an overall fractal design as in Figure 1b rather than the original sequential one. We limited our investigation to this one realization and left the study of other (possibly more complex) designs for future work. We followed the FractalNet implementation in regards to dropout where the dropout rate for a module were 0%, 10%, 20%, or 30%, depending on the depth of the module in the architecture. This choice for dropout rates were not found by experimentation and better values are possible. The local drop-path rate in the fractal-joins were fixed at 15%, which is identical to the FractalNet implementation.

Freeze-drop-path introduces four new parameters. The first is whether the active branch is chosen stochastically or deterministically. If it is chosen stochastically, a random number is generated and the active branch is assigned based on which interval it falls in (intervals will be described shortly). If it is deterministically, a parameter is set by the user as to the number of iterations in one cycle through all the branches (we called this parameter *num_iter_per_cycle*). In our Caffe implementation of the freeze-drop-path unit, the bottom input specified first is assigned as branch 1, the next is branch 2, then branch 3, etc. The next parameter indicates the proportion of iterations each branch should be active relative to all the other branches. The first type of interval uses the square of the branch number (i.e., 1, 4, 9, 16, ...) to assign the interval length for that branch to be active, which gives the more update iterations to the higher numbered branches. The next type gives the same amount of iterations to each branch. Our experiments showed that the first interval type works better (as we expected) and was used to obtain the results in Section 4.2. In addition, our experiments showed that the stochastic option works better than the deterministic option (unexpected) and was used for Section 4.2 results.

²<https://github.com/gustavla/fractalnet/tree/master/caffe>

The Stagewise Boosting Network's (SBN) architecture is the same as the FoF architecture except that branches 2 and 3 are combined with a fractal-join and then combined with branch 1 in a freeze-drop-path join. The reason for combining branches 2 and 3 came out of our first experiments; if branches 2 and 3 were separate, the performance deteriorated when branch 2 was frozen and branch 3 was active. In hindsight, this is due to the weights in the branch 2 path that are also in branch 3's path being modified by the training of branch 3. The Taylor series network has the same architecture as SBN with the addition of squaring the branch 2 and 3 combined activations before the freeze-drop-path join.

For all of our experiments, we trained for 400 epochs. Since the training used 8 GPUs and each GPU had a batchsize of 25, 400 epochs amounted to 100,000 iterations. We adopted the same learning rate as the FractalNet implementation, which started at 0.002 and dropped the learning rate by a factor of 10 at epochs 200, 300, and 350.