arXiv:1902.04698v2 [stat.ML] 15 Feb 2019

Identity Crisis: Memorization and Generalization under Extreme Overparameterization

Chiyuan Zhang chiyuan@google.com
Samy Bengio bengio@google.com
Moritz Hardt hardt@berkeley.edu
Yoram Singer singer@google.com

Abstract

We study the interplay between memorization and generalization of overparametrized networks in the extreme case of a single training example. The learning task is to predict an output which is as similar as possible to the input. We examine both fully-connected and convolutional networks that are initialized randomly and then trained to minimize the reconstruction error. The trained networks take one of the two forms: the constant function ("memorization") and the identity function ("generalization"). We show that different architectures exhibit vastly different inductive bias towards memorization and generalization. An important consequence of our study is that even in extreme cases of overparameterization, deep learning can result in proper generalization.

1 Introduction

Generalization properties of deep neural networks have attracted substantial interest in recent years due to their empirical success. A popular and effective approach is to use overly large network whose number of parameters often exceeds the number of training examples. When the training data is "natural", gradient-based training of overparameterized networks results nonetheless in the state-of-the-art performance. The cult belief is that gradient methods have implicit inductive bias towards simple solutions that generalize well. Alas, a distilled notion of inductive bias is not clearly established and understood. Numerous theoretical and empirical studies of inductive bias in deep learning have emerged in recent years (Dziugaite and Roy, 2016; Kawaguchi et al., 2017; Bartlett et al., 2017; Neyshabur et al., 2017; Liang et al., 2017; Neyshabur et al., 2018; Arora et al., 2018; Zhou et al., 2019). Unfortunately, these and other postmortem analyses do not tell us what are the root causes for the strong inductive bias. Another line of research tries to characterize sufficient conditions on the data and labels distribution, from linear separability (Brutzkus et al., 2018) to compact structures (Li and Liang, 2018), which could guarantee generalization of the trained networks. This direction, while very promising, have not yet identified structures that simple linear or nearest neighbor classifiers over *the original input space* cannot solve. The fact that deep neural networks in many applications significantly outperform these simpler models indicate that there is still a gap to fill in our understanding of deep neural networks.

On the optimization front, several recent papers (Allen-Zhu et al., 2018b; Du et al., 2018a,b; Zou et al., 2018) show that when a network is sufficiently large (e.g. the hidden units in each layer is polynomial in the input dimension and the number of training examples), under some mild assumptions, gradient methods are guaranteed to perfectly fit the training set. However, these results do *not* differentiate a model trained on the true input distribution from one trained on the same inputs albeit with random labels. While the first model might have good generalization the second model merely memorizes the training labels, thus not much light is further shed by these studies on the question of the inductive bias.

This paper embarks on an empirical study of inductive bias in deep learning by examining a restrictive setting. Concretely, we focus on learning the *identity* function in a regression setting. In doing so we are able to provide visualizations of various aspects of the learning process. We further constrain ourselves to the extreme case of learning from a *single* training example. This setting mimics extreme overparameterized regimes that were studied recently and mentioned above. The simplicity of the learning setting lets us distinguish between generalization (learning the identity map) and memorization (learning a constant map). Do large neural networks converge to the identity map or the constant map? Our experiments show that the answer is subtle and depends on the model architecture. In a broad set of experiments we highlight depth, random initialization, and different hyperparameters as relevant variables. In conclusion, our work shows that even in extreme cases of overparameterization, neural nets may not only resort to memorization, but can exhibit interesting inductive bias due to mindful architecture choices.

2 Related work

The consequences of overparameterized models in deep learning have been extensively studied in recently years, on the optimization landscape and convergence of SGD (Allen-Zhu et al., 2018b; Du et al., 2018a,b; Bassily et al., 2018; Zou et al., 2018; Oymak and Soltanolkotabi, 2018), as well as the generalization guarantees under stronger structural assumptions of the data (Li and Liang, 2018; Brutzkus et al., 2018; Allen-Zhu et al., 2018a). Another line of related work is the study of the implicit regularization effects of SGD on training overparameterized models (Neyshabur et al., 2014; Zhang et al., 2017; Soudry et al., 2018; Shah et al., 2018).

The behaviors of memorization in learning are also explicitly studied from various perspectives such as prioritizing learning of simple patterns (Arpit et al., 2017) or perfect interpolation of the training set (Belkin et al., 2018). More recently (during the writing of this paper), Radhakrishnan et al. (2018) study the effects of the downsampling operator in convolutional auto-encoders on image memorization. They use an empirical framework similar to ours, fitting ConvNets to the autoregression problem with few training examples. We focus on investigating the general inductive bias in the extreme overparameterization case, and study a broader range of network types without enforcing a bottleneck in the architectures.

3 Learning from a single example

In this paper, we focus on the learning of the identity function $\mathsf{I} : x \mapsto x$ with deep neural networks. We study the extreme overparameterized scenario with only one training example. We are interested in inspecting the mechanism of how neural networks overfit via memorization of the training data and how various kind of inductive bias come into play. Let $\hat{x}$ be the training example, we fit neural networks via the *mean squared error* loss using standard gradient descent learning setup. Various neural network architectures are studied. We explicitly ensure the configurations of the network architectures allow simple realization of the identity function (see Appendix A).

3.1 Fully connected linear networks

As a warm up, we start with the convex linear case. Consider learning the identity function with $f_W(x) = Wx$ where $x \in \mathbb{R}^d, W \in \mathbb{R}^{d\times d}$. In this case, we have a convex problem, and the optimization behaviors

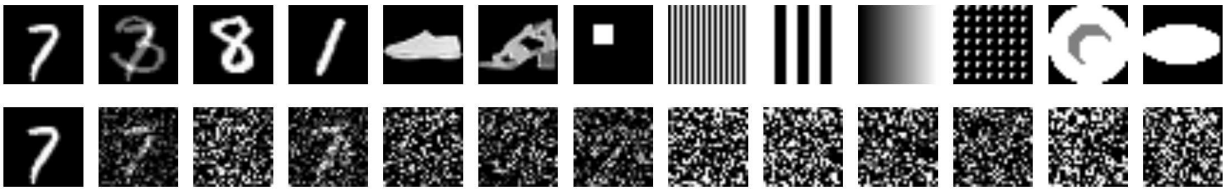

Figure 1: **Visualization of the predictions from a trained one-layer linear network.** The first row shows test inputs that consist of the (single) training digit (in the first column), linear combination of two digits, random digits from MNIST test set, random images from the Fashion-MNIST dataset, and some generated image patterns. The second row shows the corresponding predictions.

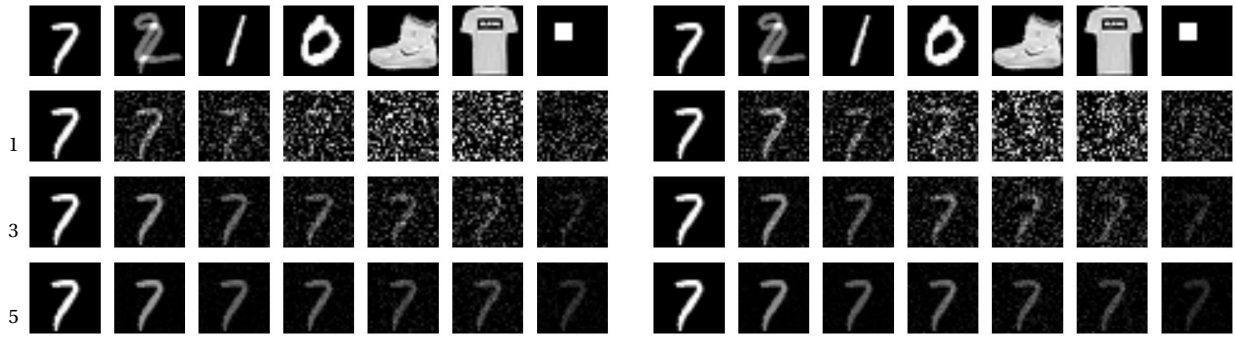

(a) Hidden dimension 784 (= input dimension)

(b) Hidden dimension 2048

Figure 2: **Visualization of predictions from trained multi-layer linear networks.** The first row shows the test images, and the remaining rows shows the prediction from a trained linear network with 1, 3, and 5 hidden layers, respectively.

are well understood. There is no unique solution to the empirical risk minimization problem due to overparameterization. But gradient descent converges to a fixed solution once the initialization is realized. In particular, let W^0 be the randomly initialized weights, then it is easy to show (see Appendix B) that the unique global minimizer exists as

$$\hat{f}_W(x) = \frac{\hat{x}^\top x}{\|\hat{x}\|_2^2} \cdot \hat{x} + W^0 \left(x - \frac{\hat{x}^\top x}{\|\hat{x}\|_2^2} \cdot \hat{x} \right) \tag{1}$$

In this case, we can fully characterize the prediction of a model trained on $\hat{x}$: the test example x is decomposed into the component in the direction of $\hat{x}$ and the orthogonal component. The component parallel to $\hat{x}$ will make the prediction look like $\hat{x}$, while the orthogonal component completely depends on the random initialization.

In this case, The learning algorithm has a strong inductive bias that converges to a unique solution once an initialization W^0 is given. But this inductive bias, unsurprisingly, is not magically leading us to generalize under this overparameterized situation. The trained model overfits as it fails to learn the identity function. More specifically, it predict well near the vicinity (measured by correlations) of the training example $\hat{x}$, but the predictions become random as it moves further away. In particular, when the test example x is orthogonal to $\hat{x}$, the prediction is completely random. Figure 1 shows an example of the predictions on various test images from a convex linear model trained on a single MNIST digit. As our calculation shows, for test patterns that resembles the training digit, the predictions are closer to $\hat{x}$, but

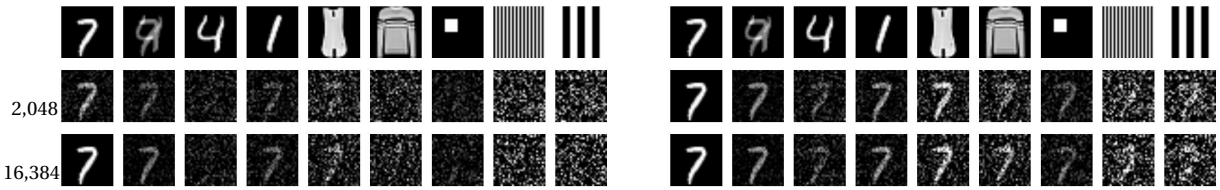

(a) Training bottom layer only

(b) Training both layers

Figure 3: **Visualization of predictions from two-layer ReLU networks.** The first row shows the test images, and the remaining rows shows the predictions from trained models with hidden dimension 2,048 and 16,384, repectively.

for unrelated test patterns, the predictions look like white noise.

Note that although the linear model does not magically learn the identity function, the overfitting behaves "nicely" as the random predictions on unfamiliar test examples can be treated as "unknown", as oppose to over-confidently predicting some wrong answer. On the other hand, Figure 2 shows the results of multi-layer *linear* networks. Because of the absence of non-linear activation functions, a multi-layer linear network (without bottleneck in the hidden dimensions) has essentially the same representation power as a single-layer linear network. But unlike the one-layer convex case, the learning dynamics is non-convex. The predictions on various test examples show that it also has very different inductive biases for the predictions on the unseen test examples. More specifically, the model with one hidden layer still resemble the convex case, but as the depth increases, the model starts to bias towards a constant function that maps everything to the single training image. The depth of the architecture shows a stronger effect on the inductive bias than the width. For example, the network with one hidden layer of dimension 2048 has 3,214,096 parameters, more than the 2,461,760 parameters of the network with three hidden layers of dimension 784. But the latter behaves more differently (from the convex case).

3.2 Two-layer fully connected networks

Li and Liang (2018) provides a theoretical characterization of the optimization and generalization of learning with a two-layer ReLU neural network. They show that when the data consist of well separated clusters (i.e. the cluster diameters are much smaller than the distances between each cluster pair), an overparameterized two-layer ReLU network can be trained and generalize on such data. To simplify the analysis, they study a special case where only the bottom layer weights are learned — the weights in the top layer are randomly initialized and fixed.

We study the problem of learning a two-layer ReLU network under our framework of learning the identity target. In Figure 3(a) and Figure 3(b), we compare the case of learning the bottom layer only and learning both layers. The visualization shows that the two cases demonstrates different inductive biases for predictions on unseen test images. In particular, when only the first layer is trained, the predictions on non-digit test examples look random; but when both layers are trained, the learned network sees the digit '7' (the image used for training) from all test images.

Note our observation is not contradicting the generalization results in Li and Liang (2018). The theoretical results require very well separated and clustered data distribution. Therefore the main concern is in the near vicinity of the training examples. While for our case, we are mostly interested in studying the (interpolation) behaviors on the test examples far away from the training set.

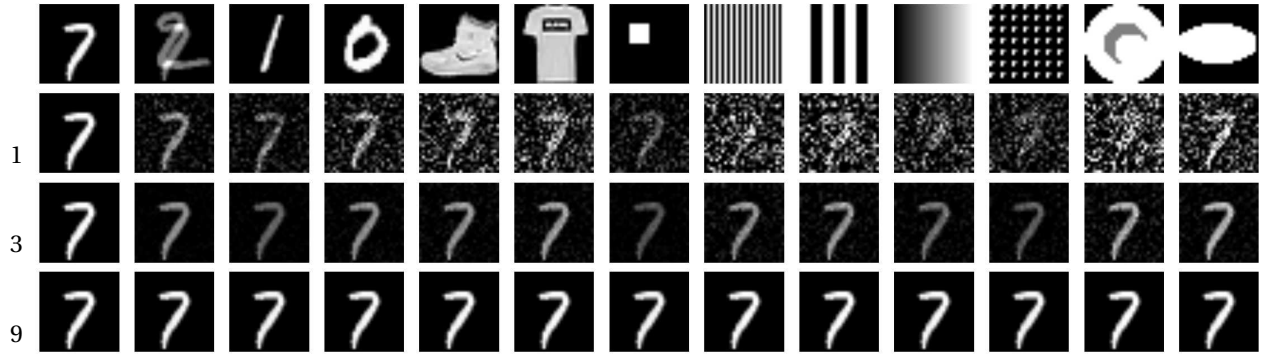

Figure 4: **Visualization of predictions from multi-layer ReLU networks.** The first row shows the test images, and the remaining rows show the predictions from trained multi-layer ReLU FCNs with 1, 3, and 9 hidden layers.

The situation when only the bottom layer is trained can be explained via a similar approach as in Section 3.1 for one-layer networks. Although we no longer have a closed form solution for the trained weights, it can be easily shown (see Appendix C) that the solution found by gradient descent is always parameterized as

$$W^T = W^0 + u^T \hat{x}^\top \tag{2}$$

where $u^T \in \mathbb{R}^m$ summarizes the efforts of gradient descent up to time T. At this point, it is easy to see if the test example is orthogonal to $\hat{x}$ (i.e. $\hat{x}^\top x = 0$), the prediction solely depends on the randomly initialized values in W^0, therefore, will be random (and can be characterized if we know the distribution for parameter initialization).

However, when both layers are trained, the upper layer weights are also tuned to make the prediction fit the training output. In particular, the learned weights in the upper layer depend on W^0. Therefore, the randomness arguments in Appendix C can no longer apply even for test examples orthogonal to $\hat{x}$. As the empirical results show, the behavior is indeed different.

3.3 Non-Linear multi-layer fully connected networks

In this section, we study the general multi-layer fully connected networks (FCNs) with the Rectified Linear Unit (ReLU) activation functions.

Figure 4 visualizes the predictions from trained ReLU FCNs with various number of hidden layers. The observation is similar to the case of multi-layer linear networks, but more pronounced. In particular, the networks are biased towards encoding the constant map with higher confidence and less prediction noise as the depth increases. To quantitatively evaluate the learned model, we measure the performance via correlation[1] to the two reference functions: the identity function, and the constant function that always maps everything to the training point $\hat{x}$. To evaluate the predictions on test images with various similarity to the training image, we generate the test images by randomly sample the pixels and then manually rotate the test vector to a given correlation $\rho \in [0,1]$ to $\hat{x}$. We also match the norm of the generated test images to $\hat{x}$. So for $\rho = 0$, the test images are orthogonal to $\hat{x}$, while for $\rho = 1$, the test images equal $\hat{x}$. The results for FCNs with different number of hidden layers are shown in Figure 5. The results for *linear* FCNs are also shown for comparison. The linear and ReLU FCNs behave similarly when measuring

[1] See Appendix I for the results in MSE.

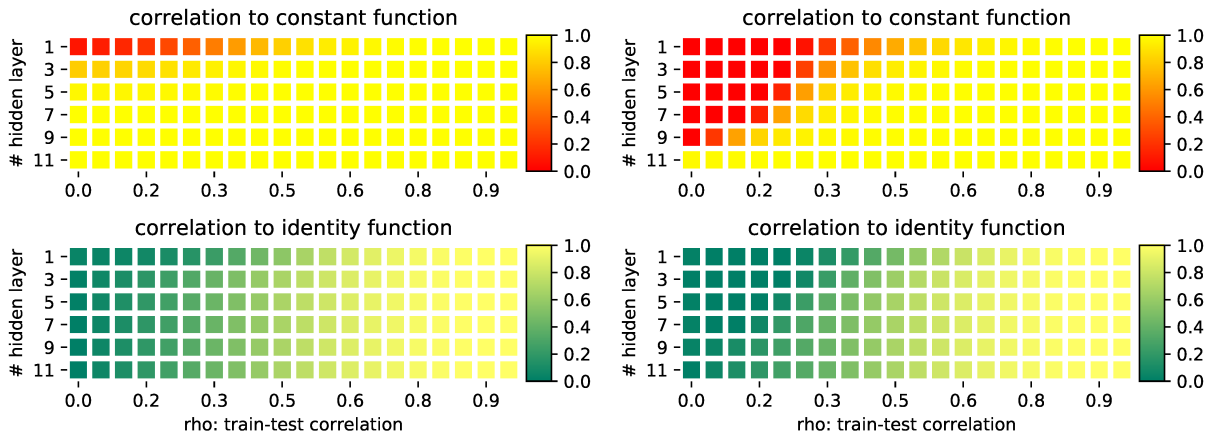

(a) ReLU FCNs

(b) Linear FCNs

Figure 5: **Quantitative evaluation of the learned model on randomly generated test samples at various angles (correlation) to the training image.** The horizontal axis shows the train-test correlation, while the vertical axis indicate the number of hidden layers for the FCNs being evaluated. The heatmap shows the similarity (measured in correlation) between the model prediction and the reference function (the constant function or the identity function). (a) shows the results for FCNs with the ReLU activation function; (b) shows the results for linear FCNs.

the correlation to the identity function: neither of them performs well for test images that are nearly orthogonal to $\hat{x}$. For the correlation to the constant function, ReLU FCNs overfit more quickly than linear FCNs when the depth increases. This is consistent with our previous visual inspections: for shallow models, the networks learn neither the constant nor the identity function, as the predictions on nearly orthogonal examples are random.

3.4 Convolutional networks

In this section, we study the inductive biases for convolutional neural networks (with ReLU activation). The prediction results on various test patterns from trained ConvNets with different depth are shown in Figure 6. Comparing to general fully connected FCNs, ConvNets have strong structural constraints that limit the receptive field of each neuron to a spatially local neighborhood, and the same weights are re-used across the spatial dimension. The two constraints indeed match the structure of the identity target function (see Appendix A.3 for an example of constructing the identity function with ConvNets).

The figure shows that except for some artifacts on the boundaries, the ConvNets with depth up to 5 learn good approximation to the identity function. Quantitative evaluation on test examples at various angles to the training image is shown in Figure 7, formatted in the same way as the heatmap for FCNs. The quantitative results are consistent with the visualizations: shallow ConvNets are able to learn the identity function from only one training example; very deep ConvNets bias towards the constant function; the ConvNets with intermediate depth correlate well with neither the identity nor the constant function. However, unlike FCNs that produce white-noise like predictions, from the visualization we see that ConvNets behave like edge detectors.

Unlike the FCNs, ConvNets preserve the spatial relation between neurons in the hidden layers, so we can easily visualize the intermediate layers as images in comparison to the inputs and outputs, to gain

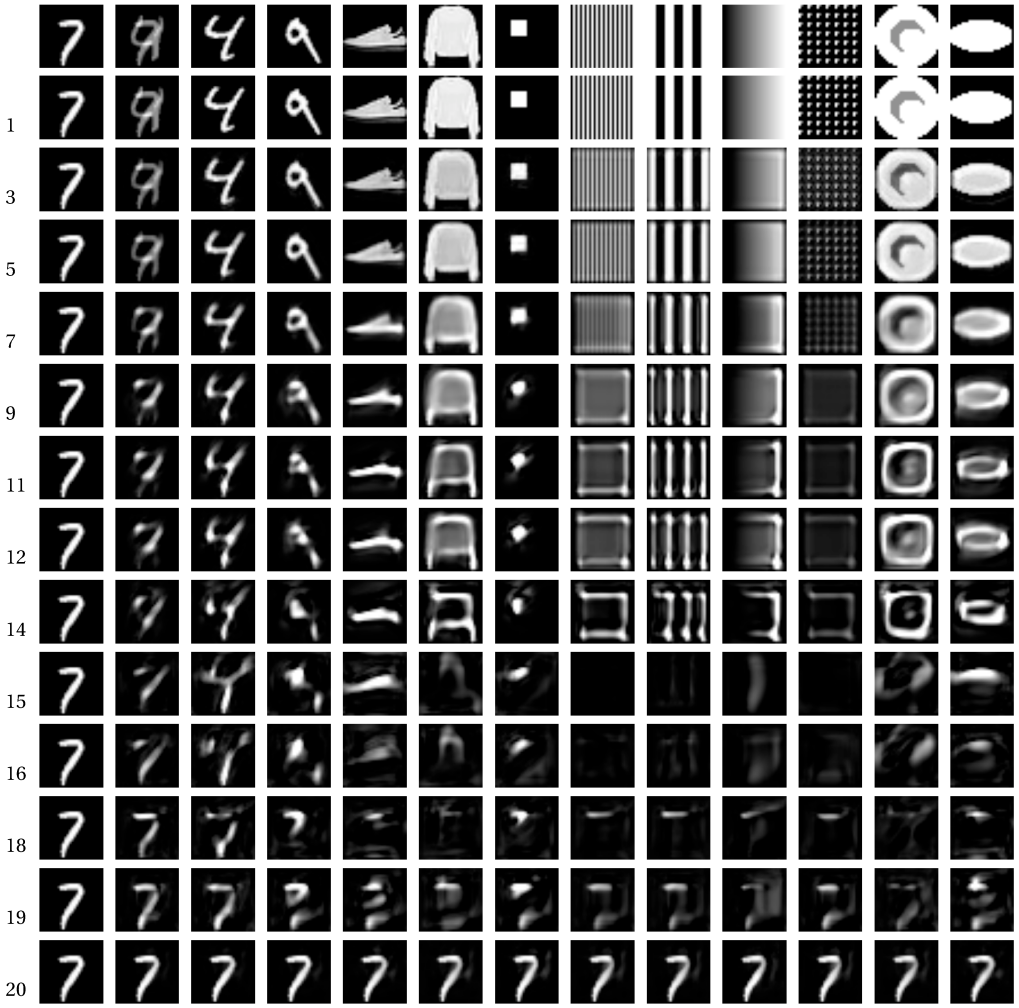

Figure 6: **Visualization of predictions from ConvNets trained with one MNIST example.** The first row shows the test images, and the remaining rows shows the predictions. The number on each row indicate the depth of the ConvNet. Each hidden layer is convolution with 5×5 kernel and 128 channels.

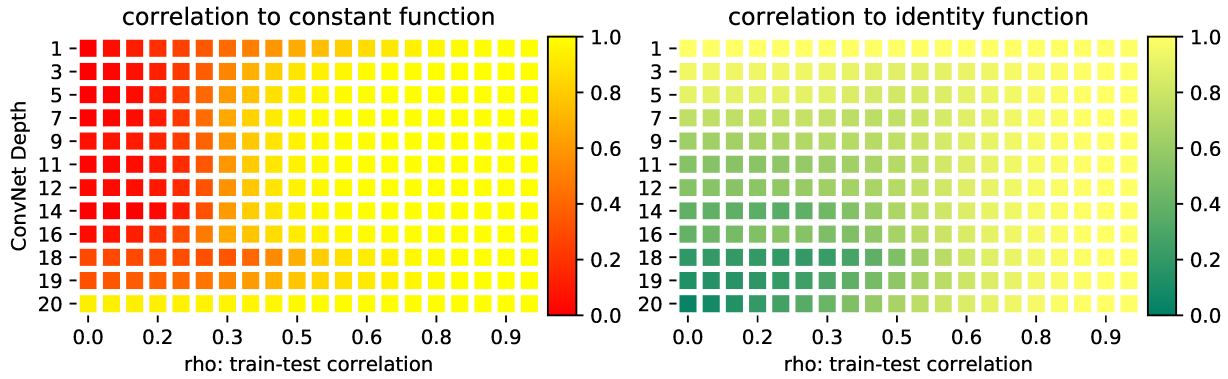

Figure 7: **Evaluation of the predictions of ConvNets on test examples at different angle to the training image.** Heatmap is formatted the same way as in Figure 5.

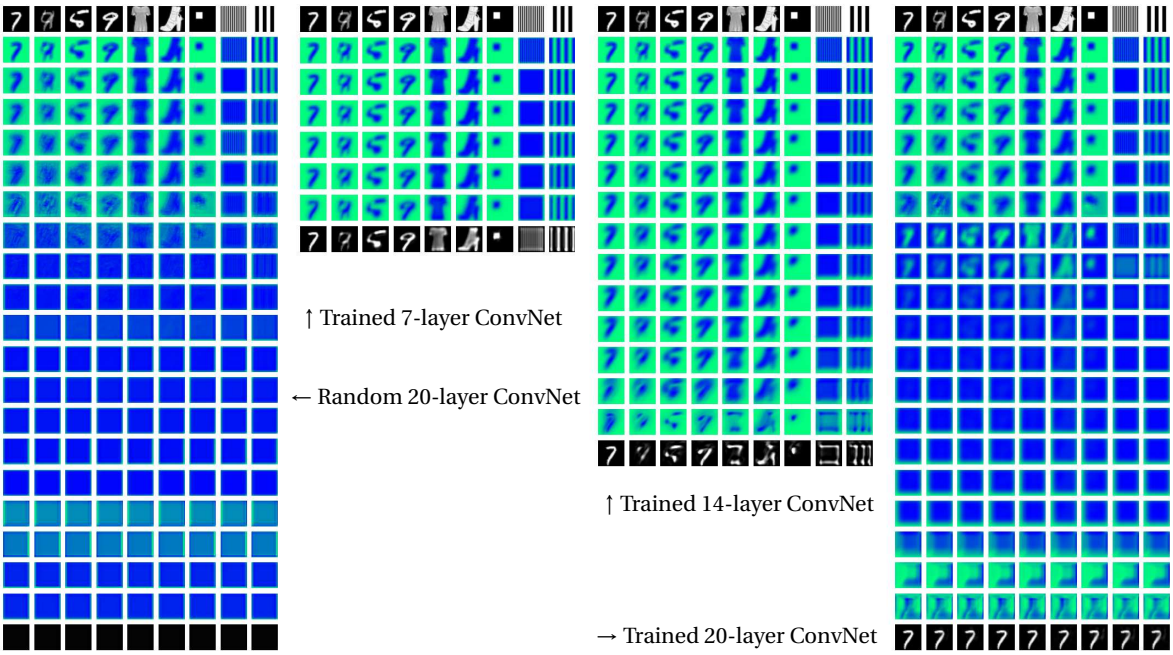

Figure 8: **Visualization the intermeidate layers of ConvNets with different number of layers.** The first column shows a randomly initialized 20-layer ConvNet (random shallower ConvNets look similar to the truncation of this). The rest of the columns show the trained ConvNets with various number of layers.

more insights on how the networks are computing the functions layer-by-layer. In Figure 8, we visualize the intermediate layer representations on some test patterns for ConvNets with different depth. In particular, for each example, the outputs from a convolutional layer in an intermediate layer is a three dimensional tensor of shape (#channel, height, width). To get a compact visualization for multiple channels in each layer, we compute SVD and visualize the top singular vector as a one-channel image. Please see Appendix E for some alternative visualizations.

In the first column, we visualize a 20-layer ConvNet at random initialization[2]. As expected, the randomly initialized convolutional layers gradually smooth out the input images. The shape of the input images are (visually) wiped out after around 8 layers of (random) convolution. On the right of the figure, we show several trained ConvNets with increasing depths. For a 7-layer ConvNet, the holistic structure of inputs are still visible all the way to the top at random initialization. After training, the network approximately renders an identity function at the output, and the intermediate activations also become less blurry. Next we show a 14-layer ConvNet, which fails to learn the identity function. However, it manages to recover meaningful information in the higher layer activations that were (visually) lost in the random initialization. On the other hand, in the last column, the network is so deep that it fails to make connection from the input to the output. Instead, the network start from scratch and constructs the digit '7' from empty and predict everything as '7'. However, note that around layer-8, we see the activations depict slightly more clear structures than the randomly initialized network. This suggests that some efforts have been made during the learning, as opposed to the case that the bottom layers not being learned due to complete gradient vanishing. Please refer to Appendix D for further details related to potential gradient vanishing problems.

[2]Shallower ConvNets at random initialization can be well represented by looking at a (top) subset of the visualization.

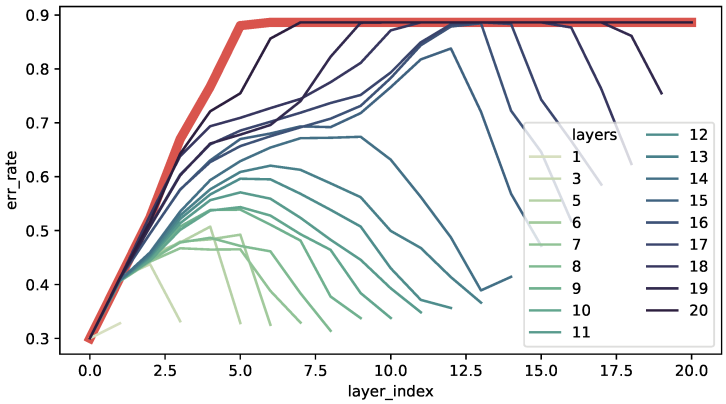

Figure 9: **Measure of the representation collapsing at each layer for trained ConvNets with different depth.** The error rate is measured by feeding the representations computed at each layer to a simple averaging based classifier on the MNIST test set. The error rate at each layer is plotted for a number of trained ConvNets with different depth. The thick semi-transparent red line shows the curve for an *untrained* 20-layer ConvNet for reference.

To get a quantitative evaluation of how much information is lost in the intermediate layer, we use the following simple criterion to measure the layerwise representations. Take the MNIST dataset, for each layer in the network (trained on a single example), collect the representations obtained by feeding each item in the dataset through the network up to that layer, then do a simple similarity based global averaging classification and measure the error rate. Specifically, the prediction for each test example is the argmax of the mean vector of the (one-hot) training labels, weighted by the correlation between the test example and each training example, computed by the representation from the layer we want to inspect.

This metric does *not* quantify how much information is preserved as the representation propagate through layers in the information theoretical sense, as the information could still be present but encoded in complicated ways that makes it hard for simple averaging-based classifier to pick out the signals. But it provides a simple metric for our case: using the raw representation as the baseline, if a layer represents the identity function, then the representation at that layer will have similar error rate to the raw representation; on the other hand, if the a layer collapses to the identity function, then the corresponding representation will have error rate close to random. The results are plotted in Figure 9. The error rate curve for a randomly initialized 20-layer ConvNet is also shown as reference: at random initialization, the smoothing effect makes the representations beyond around layer 5-6 almost useless for our simple classifier. After training, as the error rates for the output layers decrease, the curves generally form "concave" patterns. This is consistent with the visualizations in Figure 8: the trained networks try to recover the smoothed out intermediate layer representations and make connections between the inputs and the outputs. But if the gap is too big to push all the necessary information through, the network will try to infer the input-output relation using partial information, resulting in models that behave like edge detectors. Finally, for the case of 20 layers, the curve shows that the bottom few layers do get small improvements in error rate, but the big gap between inputs and outputs drives the network to learn the constant function instead.

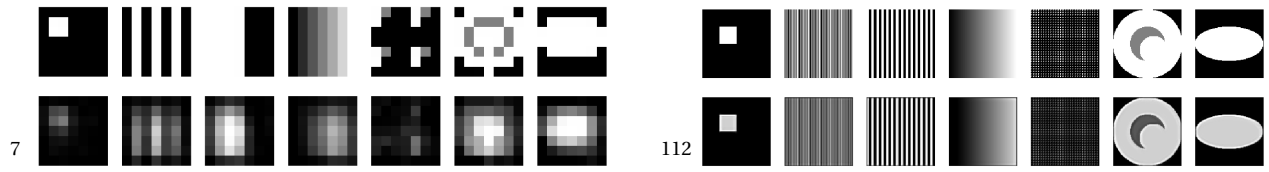

Figure 10: **Visualization of a 5-layer ConvNet on test images of different sizes**. The two subfigures show the results on 7×7 inputs and 112×112 inputs, respectively.

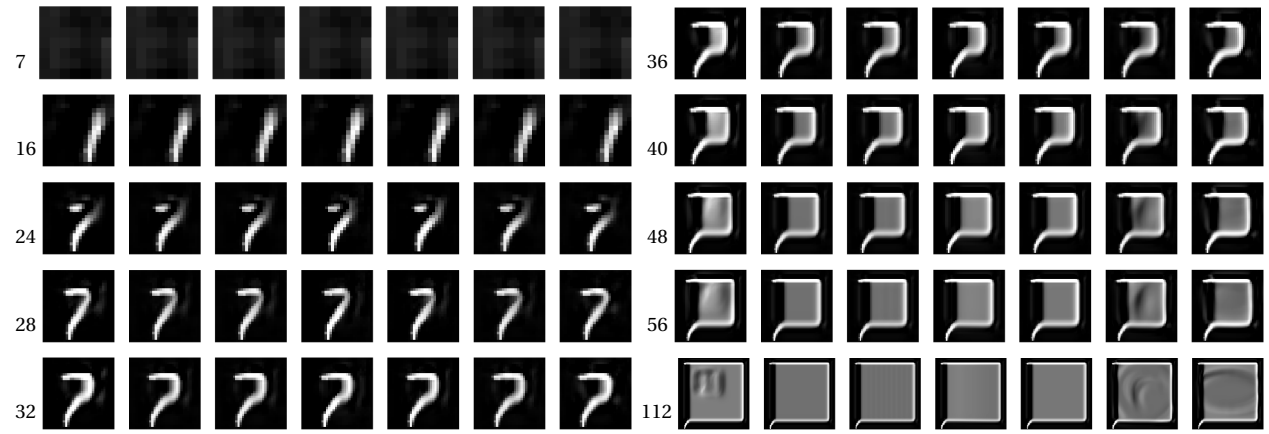

Figure 11: **Visualization of the predictions of a 20-layer ConvNet on test images of different sizes (indicated by the number on each row).** The input patterns are the same as in Figure 10 (constructed in different resolutions), which are not shown for brevity.

3.5 Robustness of inductive biases

Being able to learn the identity function from a single example shows a strong inductive bias of (not too deep) ConvNets. On the other hand, learning the constant function that maps everything to the same output is also a strong inductive bias, as the training objective never explicitly asked the model to do so. The semantics of the spatial locations of neurons in ConvNets allow us to do some extra analysis to investigate the encoding process of the learned function and the robustness of the inductive biases.

Variable input image sizes. The ConvNets are naturally agnostic to the input image sizes. We can apply a trained network to a variety of different input sizes. Figure 10 visualizes the predictions of a trained 5-layer ConvNets on very small 7×7 and large 112×112 input images. We found that the learned identity map generally holds up against larger (than the training) input sizes (see Appendix F for a more complete set of results). However, on small inputs, the predictions no longer match well with the inputs. Note ConvNets are in principle capable of encoding the identity function for arbitrary inputs and filter sizes, for example, via the construction in Appendix A.3.

Figure 11 shows the predictions on the same set of input patterns (of different resolutions) by a 20-layer ConvNet which learns the constant function on a 28×28 image. We found that the learned constant map holds up in a smaller range of input sizes (than the learned identity map). But it is interesting to see smooth changes as the input sizes increases to show the network's own notion of "7".

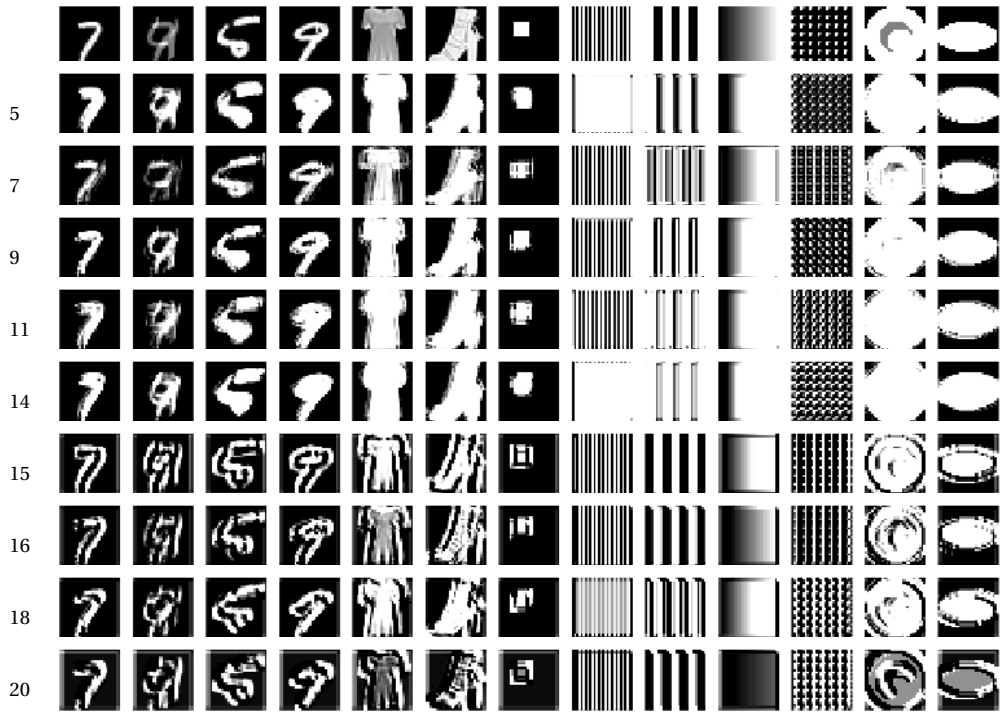

Figure 12: **Visualizing only the final layer in trained networks.** The first row are the input images, which are directly fed into the final layer of trained networks (skipping the bottom layers). The remaining rows shows the predictions from the top layers of ConvNets, with the numbers on the left indicating their (original) depth.

The upper subnetwork. In the visualization of intermediate layers (Figure 8), the intermediate layers actually represent the "lower" subnetwork from the inputs. Here we investigate the "upper" subnetwork. Thanks again to the spatial structure of ConvNets, we can skip the lower layers and feed the test patterns directly to the intermediate layers and still get interpretable visualizations[3]. Figure 12 shows the results for the top-one layer from ConvNets with various depths. A clear distinction can be found at 15-layer ConvNet, which according to Figure 6 is where the networks start to bias away from edge detector and towards the constant function. See Appendix G for visualization of larger chunk of upper subnetworks.

3.6 Varying other factors

Studies on how other factors of the ConvNet architecture affect the inductive biases are presented here. More detailed results and discussions on this are presented in Appendix H.

Figure 13 shows the mean correlation to the constant and the identity function across the network depths when we train with different input (and output) image sizes. The training examples are resized versions of the same image. We can see that with smaller training images, the ConvNets are more easily biased towards the constant function as the depth increases. Meanwhile, it is also less biased towards the identity function when trained with smaller images.

[3] Specifically, the intermediate layers expect inputs with multiple channels, so we repeat the grayscale inputs across channels to match the expected input shape.

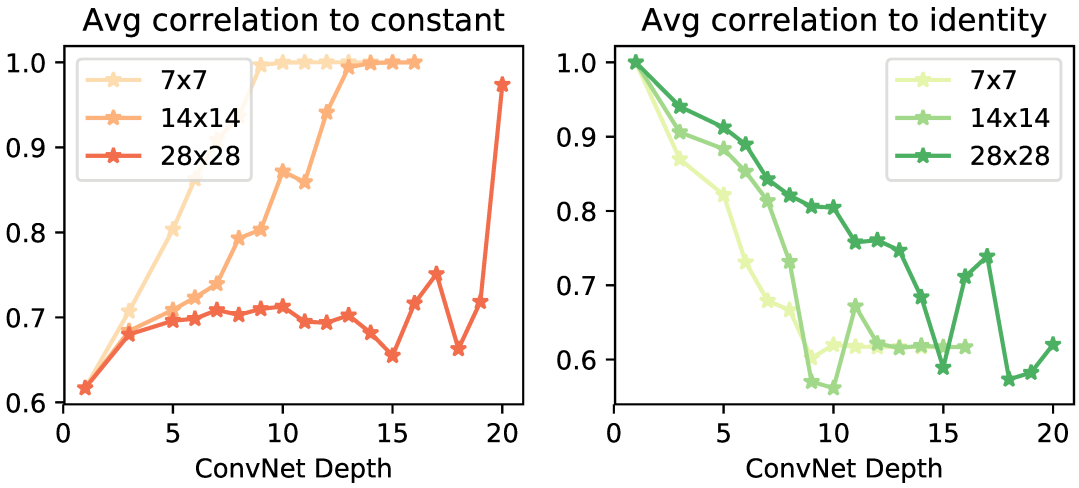

Figure 13: **Comparing bias towards constant and identity when trained with different image sizes.** The x-axis is the depth of the ConvNets, while the y-axis is the mean correlation (average of each row from the heatmaps like in Figure 7). Each curve corresponds to training with a different image size.

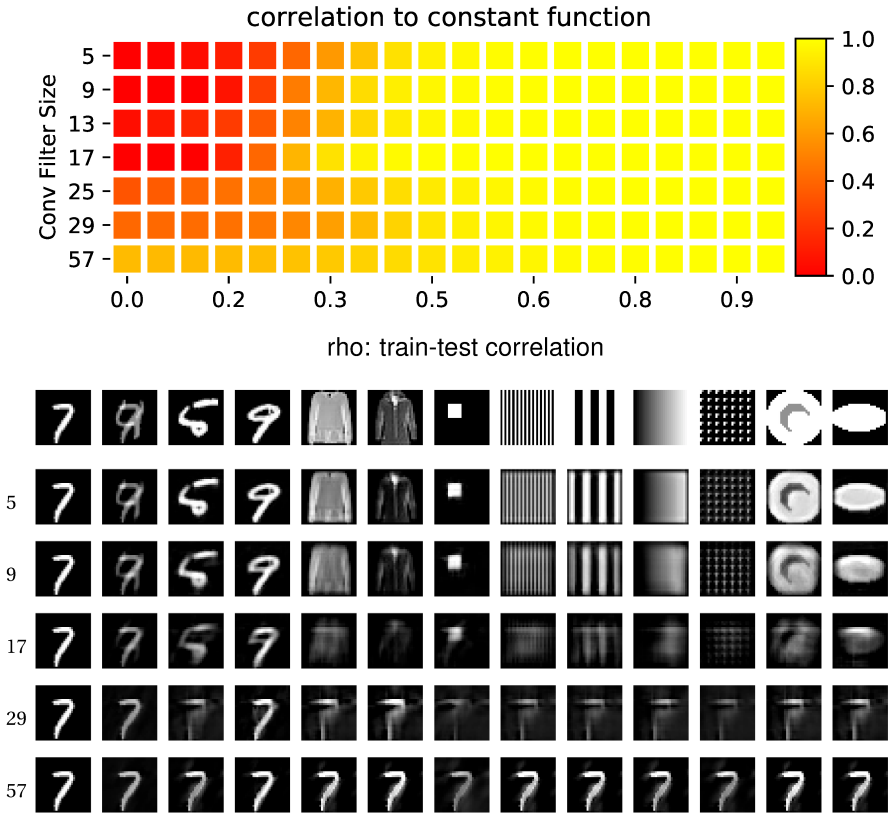

Figure 14: **Inductive bias of a 5-layer ConvNet with varying convolutional filter size.** The heatmap is arranged in the same way as Figure 7, except that the rows correspond to ConvNets with different filter sizes. On the bottom we visualize the predictions from ConvNets with a few selected filter sizes: the first row is the inputs, and the remaining rows are predictions, with the numbers on the left indicating the corresponding filter sizes.

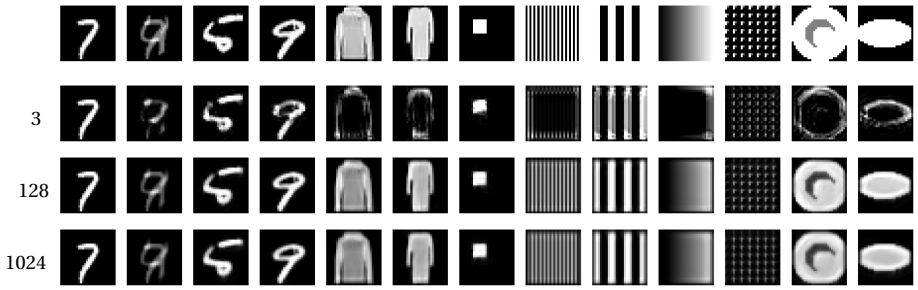

Figure 15: **Visualization of the predictions from 5-layer ConvNets with various number of hidden channels.** Numbers on the left indicate the number of channels. For the intermediate layers, the numbers indicate the number of both the input and the output channels. The number of input channels for the bottom layer and the output channels for the top layer are decided by the data (one for our case).

Figure 14 illustrates the inductive bias towards the constant function with varying convolutional filter size from 5×5 to 57×57. The heatmap shows that varying the filter sizes does not strongly affect the bias towards the constant function until very large values. The visualization of the predictions indicate that moderate sized filter sizes make the predictions more blurry, but the holistic structures of the inputs are still preserved. With extremely large filter sizes that cover the whole spatial domain in the inputs, the behavior of ConvNets starts to resemble FCNs.

Figure 15 show the results when the number of channels in convolution changes. Note the construction in Appendix A.3 shows that two channels are enough to encode the identity function for grayscale inputs. However, the results with three channels are missing a lot of center contents in many predicted images. This is not underfitting since the network reconstructs the training image (first column) correctly. On the other hand, the aggressively over-parameterized network with 1024 channels (~25M parameters per middle-layer convolution module) does not seem to suffer from overfitting.

4 Conclusions

We presented empirical studies of the extreme case of overparameterization when learning from a single example. We investigated the interplay between memorization and generalization in deep neural networks. By restricting the learning task to the identity function, we sidestepped issues such as the underlying optimal Bayes error of the problem and the approximation error of the hypothesis classes. This choice also facilitated rich visualization and intuitive interpretation of the trained models. Under this setup, we investigate gradient-based learning procedures with explicit memorization-generalization characterization. Our results indicate that different architectures exhibit vastly different inductive bias towards memorization and generalization.

Acknowledgments The authors would like to thank Mike Mozer, Kunal Talwar, Hanie Sedghi and Rong Ge for helpful discussions.

References

Allen-Zhu, Z., Li, Y., and Liang, Y. (2018a). Learning and generalization in overparameterized neural networks, going beyond two layers. *CoRR*, 1811.04918.

Allen-Zhu, Z., Li, Y., and Song, Z. (2018b). A convergence theory for deep learning via Over-Parameterization. *CoRR*, arXiv:1811.03962.

Arora, S., Ge, R., Neyshabur, B., and Zhang, Y. (2018). Stronger generalization bounds for deep nets via a compression approach. *CoRR*, arXiv:1802.05296.

Arpit, D., Jastrzębski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., et al. (2017). A closer look at memorization in deep networks. In *International Conference on Machine Learning*.

Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. (2017). Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6240–6249.

Bassily, R., Belkin, M., and Ma, S. (2018). On exponential convergence of SGD in non-convex over-parametrized learning. *CoRR*, arXiv:1811.02564.

Belkin, M., Hsu, D., and Mitra, P. (2018). Overfitting or perfect fitting? risk bounds for classification and regression rules that interpolate. In *Advances in Neural Information Processing Systems*.

Brutzkus, A., Globerson, A., Malach, E., and Shalev-Shwartz, S. (2018). SGD learns over-parameterized networks that provably generalize on linearly separable data. In *ICLR*.

Du, S. S., Lee, J. D., Li, H., Wang, L., and Zhai, X. (2018a). Gradient descent finds global minima of deep neural networks. *CoRR*, arXiv:1811.03804.

Du, S. S., Zhai, X., Poczos, B., and Singh, A. (2018b). Gradient descent provably optimizes over-parameterized neural networks. *CoRR*, arXiv:1810.02054.

Dziugaite, G. K. and Roy, D. M. (2016). Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*.

Kawaguchi, K., Kaelbling, L. P., and Bengio, Y. (2017). Generalization in deep learning. *CoRR*, arXiv:1710.05468.

Li, Y. and Liang, Y. (2018). Learning overparameterized neural networks via stochastic gradient descent on structured data. In *NIPS*.

Liang, T., Poggio, T., Rakhlin, A., and Stokes, J. (2017). Fisher-rao metric, geometry, and complexity of neural networks. *CoRR*, arXiv:1711.01530.

Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. (2017). Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956.

Neyshabur, B., Bhojanapalli, S., and Srebro, N. (2018). A PAC-Bayesian approach to Spectrally-Normalized margin bounds for neural networks. In *ICLR*.

Neyshabur, B., Tomioka, R., and Srebro, N. (2014). In search of the real inductive bias: On the role of implicit regularization in deep learning. *CoRR*, arXiv:1412.6614.

Oymak, S. and Soltanolkotabi, M. (2018). Overparameterized nonlinear learning: Gradient descent takes the shortest path? *CoRR*, arXiv:1812.10004.

Radhakrishnan, A., Belkin, M., and Uhler, C. (2018). Downsampling leads to image memorization in convolutional autoencoders. *CoRR*, arXiv:1810.10333.

Shah, V., Kyrillidis, A., and Sanghavi, S. (2018). Minimum norm solutions do not always generalize well for over-parameterized problems. *CoRR*, arXiv:1811.07055.

Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. (2018). The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70).

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *International Conference on Learning Representations*.

Zhou, W., Veitch, V., Austern, M., Adams, R. P., and Orbanz, P. (2019). Non-vacuous generalization bounds at the ImageNet scale: a PAC-Bayesian compression approach. In *ICLR*.

Zou, D., Cao, Y., Zhou, D., and Gu, Q. (2018). Stochastic gradient descent optimizes over-parameterized deep ReLU networks. *CoRR*, arXiv:1811.08888.

A Representation of the identity function using deep networks

In this section, we provide explicit constructions on how common types of neural networks can represent the identity function. Those constructions are only proof for that the models in our study have the capacity to represent the target function. There are many different ways to construct the identity map for each network architecture, but we try to provide the most straightforward and explicit constructions. However, during our experiments, even when the SGD learns (approximately) the identity function, there is no evidence suggesting that it is encoding the functions in similar ways as described here. We put some mild constraints (e.g. no "bottleneck" in the hidden dimensions) to allow more straightforward realization of the identity function, but this by no means asserts that networks violating those constraints cannot encode the identity function.

A.1 Linear models

For a one-layer linear network $f(x) = Wx$, where $W \in \mathbb{R}^{d\times d}$, setting W to the identity matrix will realize the identity function. For a multi-layer linear network $f(x) = (\prod_\ell W_\ell)x$, we need to require that all the hidden dimensions are not smaller than the input dimension. In this case, a simple concrete construction is to set each W_ℓ to an identity matrix.

A.2 Multi-layer ReLU networks

The ReLU activation function $\sigma(\cdot) = \max(0,\cdot)$ discards all the negative values. There are many ways one can encode the negative values and recover it after ReLU. We provide a simple approach that uses hidden dimensions twice the input dimension. Consider a ReLU network with one hidden layer $f(x) = W_2\sigma(W_1x)$, where $W_2 \in \mathbb{R}^{d\times 2d}, W_1 \in \mathbb{R}^{2d\times d}$. The idea is to store the positive and negative part of x separately, and then re-construct. This can be achieved by setting

$$W_1 = \begin{pmatrix} I_d \\ -I_d \end{pmatrix}, \quad W_2 = \begin{pmatrix} I_d & -I_d \end{pmatrix}$$

where I_d is the d-dimensional identity matrix. For the case of more than two layers, we can use the bottom layer to split the positive and negative part, and the top layer to merge them back. All the intermediate layers can be set to $2d$-dimensional identity matrix. Since the bottom layer encode all the responsives in non-negative values, the ReLU in the middle layers will pass through.

A.3 Convolutional networks

In particular, we consider 2D convolutional networks for data with the structure of multi-channel images. A mini-batch of data is usually formatted as a four-dimensional tensor of the shape $B \times C \times H \times W$, where B is the batch size, C the number of channels (e.g. RGB or feature channels for intermediate layer representations), H and W are image height and width, respectively. A convolutional layer (ignoring the bias term) is parameterized with another four-dimensional tensor of the shape $\bar{C} \times C \times K_H \times K_W$, where $\bar{C}$ is the number of output feature channels, K_H and K_W are convolutional kernel height and width, respectively. The convolutional kernel is applied at local $K_H \times K_W$ patches of the input tensor, with optional padding and striding.

For one convolution layer to represent the identity function, we can use only the center slice of the kernel tensor and set all the other values to zero. Note it is very rare to use even numbers as kernel size, in which case the "center" of the kernel tensor is not well defined. When the kernel size is odd, we can set

$$W_{\bar{c}chw} = \begin{cases} 1 & \bar{c} = c,\ h = \lfloor K_H/2 \rfloor,\ w = \lfloor K_W/2 \rfloor \\ 0 & \text{otherwise} \end{cases}$$

By using only the center of the kernel, we essentially simulate a 1×1 convolution, and encode a local identity function for each (multi-channel) pixel.

For multi-layer convolutional networks with ReLU activation functions, the same idea as in multi-layer fully-connected networks can be applied. Specifically, we ask for twice as many channels as the input channels for the hidden layers. At the bottom layer, separately the positive and negative part of the inputs, and reconstruct them at the top layer.

B Closed form solution for single-layer overparameterized network

In (1), a closed form solution is provided for the global minimizer of a one-layer neural network trained on a single example. The derivation of the solution is presented here. Let $\hat{x}$ be the training example, the gradient of the empirical risk is

$$\frac{\partial \hat{R}}{\partial W} = (W - I)\hat{x}\hat{x}^\top \tag{3}$$

Gradient descent with step sizes η_t and initialization weights W^0 update weights as

$$\begin{aligned} W^T &= W^0 - \sum_{t=1}^{T} \eta_t (W^{t-1} - I)\hat{x}\hat{x}^\top \\ &:= W^0 + u_T \hat{x}^\top \end{aligned}$$

where $u_T \in \mathbb{R}^d$ is a vector decided via the accumulation in the optimization trajectory. Because of the form of the gradient, it is easy to see the solution found by gradient descent will always have such parameterization structure. Moreover, under this parameterization, a unique minimizer exists that solves the equation

$$\hat{x} = f_W(\hat{x}) = W^0\hat{x} + u\hat{x}^\top\hat{x}$$

via

$$\hat{u} = \frac{(I - W^0)\hat{x}}{\|\hat{x}\|_2^2} \tag{4}$$

Therefore, the global minimizer can be written as in (1), copied here for convenience:

$$\begin{aligned} \hat{f}_W(x) &= (W^0 + \hat{u}\hat{x}^\top)x \\ &= \frac{\hat{x}^\top x}{\|\hat{x}\|_2^2} \cdot \hat{x} + W^0\left(x - \frac{\hat{x}^\top x}{\|\hat{x}\|_2^2} \cdot \hat{x}\right) \end{aligned}$$

For the one-layer network case, the optimization problem is convex. Under standard conditions in convex optimization, gradient descent will converge to the global minimizer shown above.

The calculation can be easily extended to arbitrary target function other than the identity, as well as the case with multiple training examples. It will result in a decomposition of the test example into the

subspace spanned by the training samples, and the orthogonal subspace. In particular, when the training examples have full rank — spanning the whole input space, the model will correctly learn the identity function.

C Characterization of solution when learning only the bottom layer

In section 3.2, (2) provides a characterization of the solution for two-layer neural networks when only the bottom layer is trained. The derivation of the characterization is presented here. Let us denote

$$f_W(x) = \langle \alpha, \mathrm{ReLU}(z) \rangle, \quad z = Wx \tag{5}$$

where $W \in \mathbb{R}^{m \times d}$ is the learnable weight matrix, and $\alpha \in \mathbb{R}^m$ is randomly initialized and fixed. Let $\hat{x}$ be the training example, the gradient of the empirical risk $\hat{R}$ with respect to each row $W_{:r}$ of the learnable weight is

$$\frac{\partial \hat{R}}{\partial W_{:r}} = \frac{\partial \hat{R}}{\partial z_r} \frac{\partial z_r}{\partial W_{:r}} = \frac{\partial \hat{R}}{\partial z_r} \hat{x}^\top \tag{6}$$

Putting it together, the full gradient is

$$\frac{\partial \hat{R}}{\partial W} = \frac{\partial \hat{R}}{\partial z} \cdot \hat{x}^\top \tag{7}$$

Again, since the gradient lives in the span of the training example $\hat{x}$, the solution found by gradient descent is always parameterized as (2), which we copy here:

$$W^T = W^0 + u^T \hat{x}^\top$$

where $u^T \in \mathbb{R}^m$ summarizes the efforts of gradient descent up to time T. The same arguments applies to multi-layer neural networks. The prediction on any test example that is orthogonal to $\hat{x}$ will depend only on randomly initialized W^0 and upper layer weights. When only the bottom layer is trained, the upper layer weights will also be independent from the data, therefore the prediction is completely random. However, when all the layers are jointly trained, the arguments no longer apply. The empirical results presented later in the paper that multi-layer networks bias towards the constant function verify this.

D Measuring the change in weights of layers post training

In this section, we study the connection between the inductive bias of learning the constant function and the potential gradient vanishing problem. Instead of measuring the norm of gradient during training, we use a simple proxy that directly compute the distance of the weight tensor before and after training. In particular, for each weight tensor W^0 and initialization and $W^\star$ after training, we compute the relative ℓ_2 distance as

$$d(W^0, W^\star) := \frac{\|W^0 - W^\star\|_2}{\|W^0\|_2}$$

The results for ConvNets with various depths are plotted in Figure 16. As a general pattern, we do see that as the network architecture gets deeper, the distances at lower layers do become smaller. But they are still non-zero, which is consistent with the visualization in Figure 8 showing that even for the 20-layer ConvNet, where the output layer fits to the constant function, the lower layers does get enough updates to allow them to be visually distinguished from the random initialization.

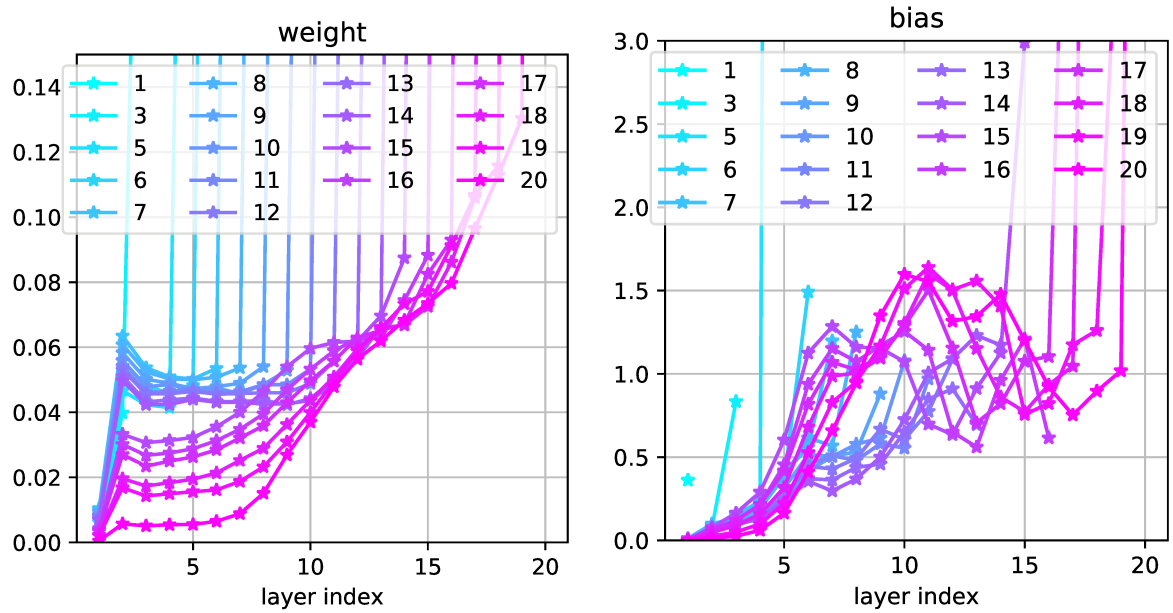

Figure 16: **The relative ℓ_2 distance of the weight and bias tensors before and after training at each layer.** The curves compare ConvNets at different depth. Most of the networks have significantly larger distances on the top-most layer. To see a better resolution at the bottom layers, we cut off the top layer in the figures by manually restricting the y axis.

In Figure 17 and Figure 18, we show the same plots for linear FCNs and FCNs with ReLU activation, respectively. We see that especially for ReLU FCN with 11 hidden layers, the distances for the weight tensors at the lower 5 layers are near zero. However, recall from Figure 4 in Section 3.3, the ReLU FCNs start to bias towards the constant function with only three hidden layers, which are by no means suffering from vanishing gradients as the plots here demonstrate.

E Alternative visualizations of the intermediate layers of ConvNets

In Section 3.4, we visualize the intermediate representations of ConvNets by showing the top singular vector across channels in each layer. We provide two alternative visualizations here showing the channel that is maximally correlated with the input image, and a random channel (channel 0). Figure 19, Figure 20 and Figure 21 illustrate a 7-layer ConvNet, a 14-layer ConvNet and a 20-layer ConvNet, respectively.

F Full results for inputs of different sizes

Figure 10 in Section 3.5 visualize the predictions of a 5-layer ConvNet on inputs of the extreme sizes of 7×7 and 112×112. Here in Figure 22 we provide more results on some intermediate image sizes.

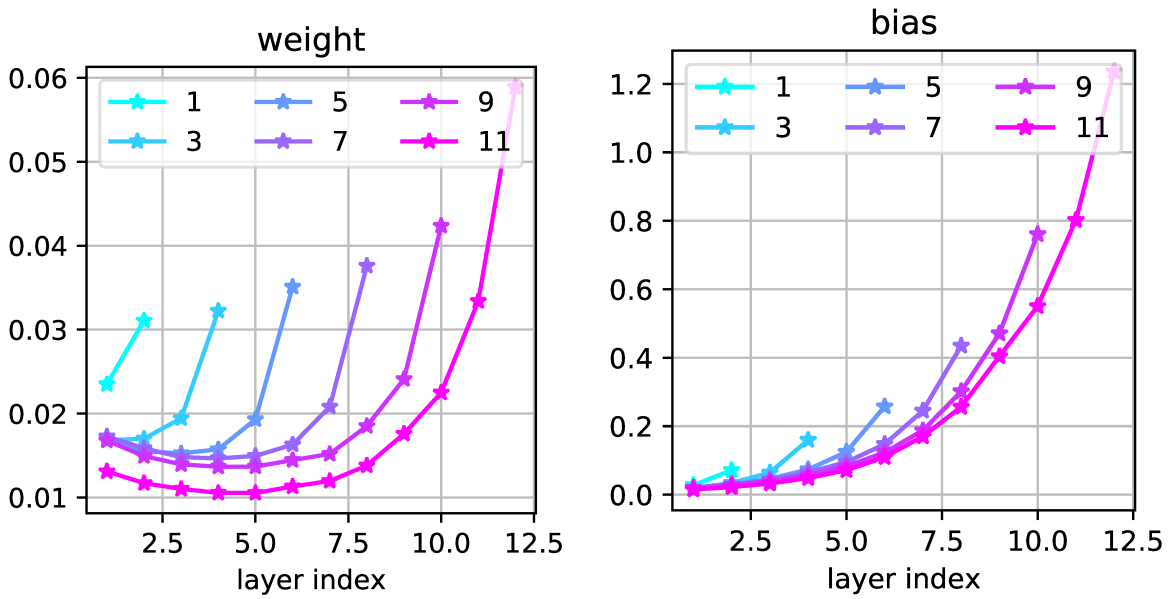

Figure 17: **The relative ℓ_2 distance of the weight and bias tensors before and after training at each layer.** The curves compare linear fully connected networks with different number of hidden layers.

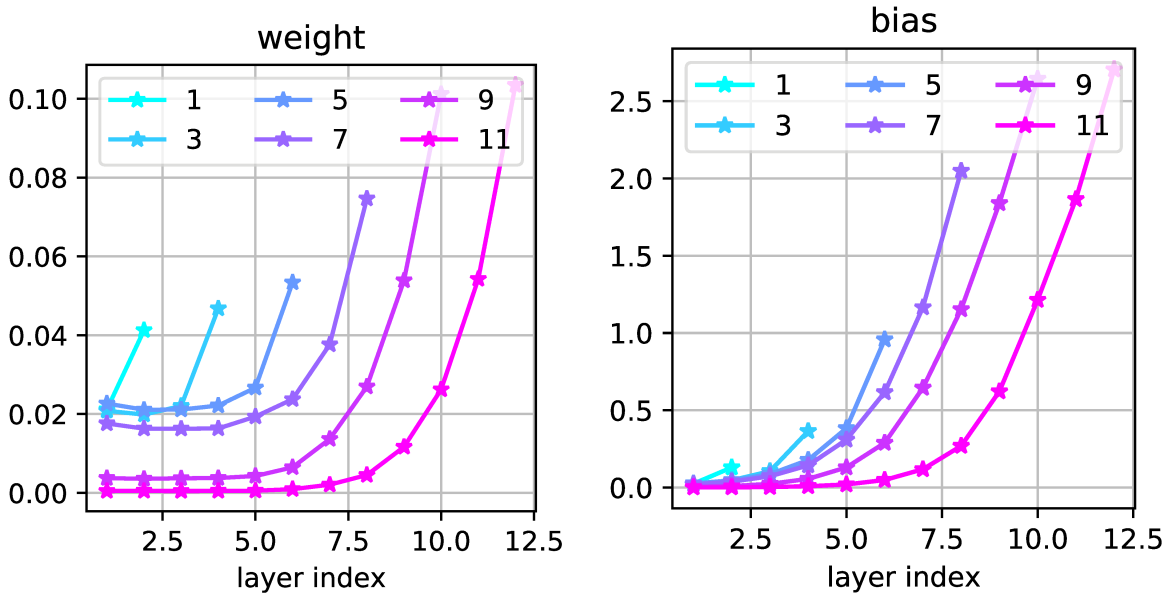

Figure 18: **The relative ℓ_2 distance of the weight and bias tensors before and after training at each layer.** The curves compare fully connected networks with ReLU activation with different number of hidden layers.

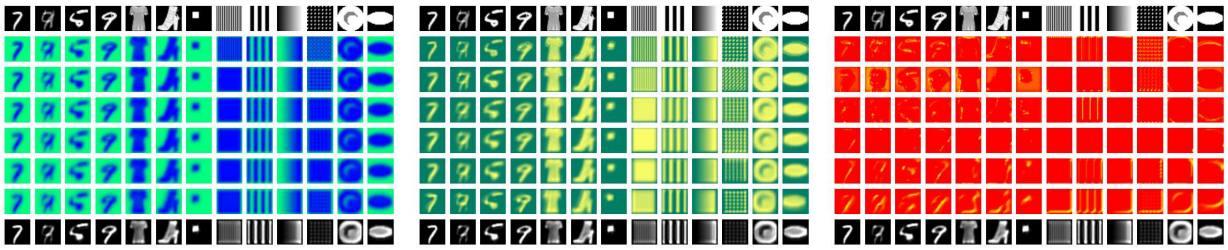

Figure 19: **Visualizing the intermediate layers of a trained 7-layer ConvNet.** The three subfigures show for each layer: 1) the top singular vector across the channels; 2) the channel that maximally correlate with the input image; 2) a random channel, respectively.

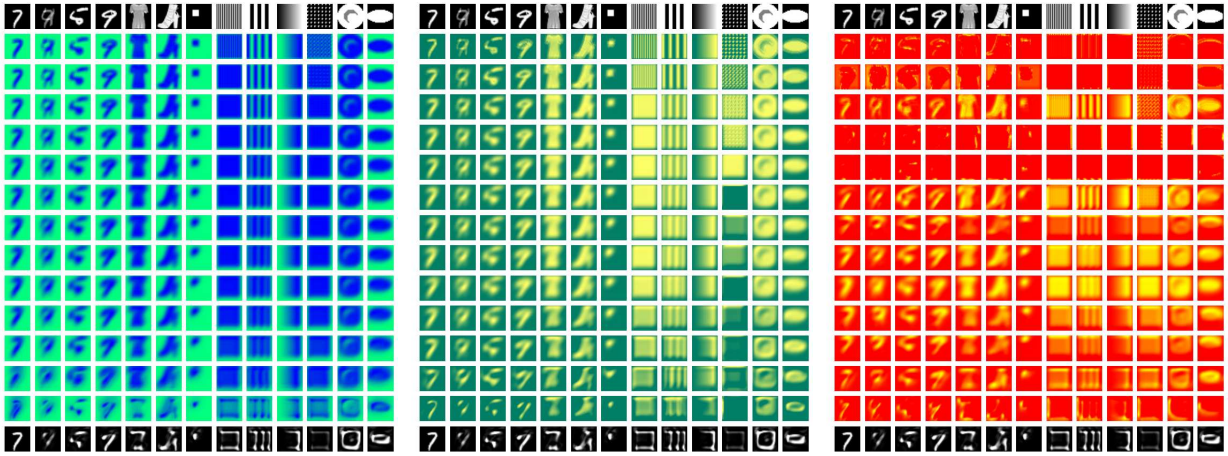

Figure 20: **Visualizing the intermediate layers of a trained 14-layer ConvNet.** The three subfigures show for each layer: 1) the top singular vector across the channels; 2) the channel that maximally correlate with the input image; 2) a random channel, respectively.

G Visualization of the upper sub-network

Figure 12 illustrated the predictions of the final layer from various trained network by directly feeding the inputs (skipping the lower layers). Further results are presented in this section. The predictions from the final two layers of each network are visualized in Figure 23. Figure 24 focuses on the 20-layer ConvNet that learns the constant map, and visualize the upper 3 layers, 6 layers and 10 layers, respectively. In particular, the last visualization shows that the 20-layer ConvNet is already starting to construct the digit "7" from nowhere when using only the upper half of the model.

H Results for further factors in ConvNets

Section 3.6 studies how the factors like the training image sizes, the convolutional filter sizes and the number of convolution channels affect the inductive bias of ConvNets. Results not included in the main text due to space limit are presented here.

Figure 25 and Figure 26 complete Figure 14 in Section 3.6 with full results that compare the inductive

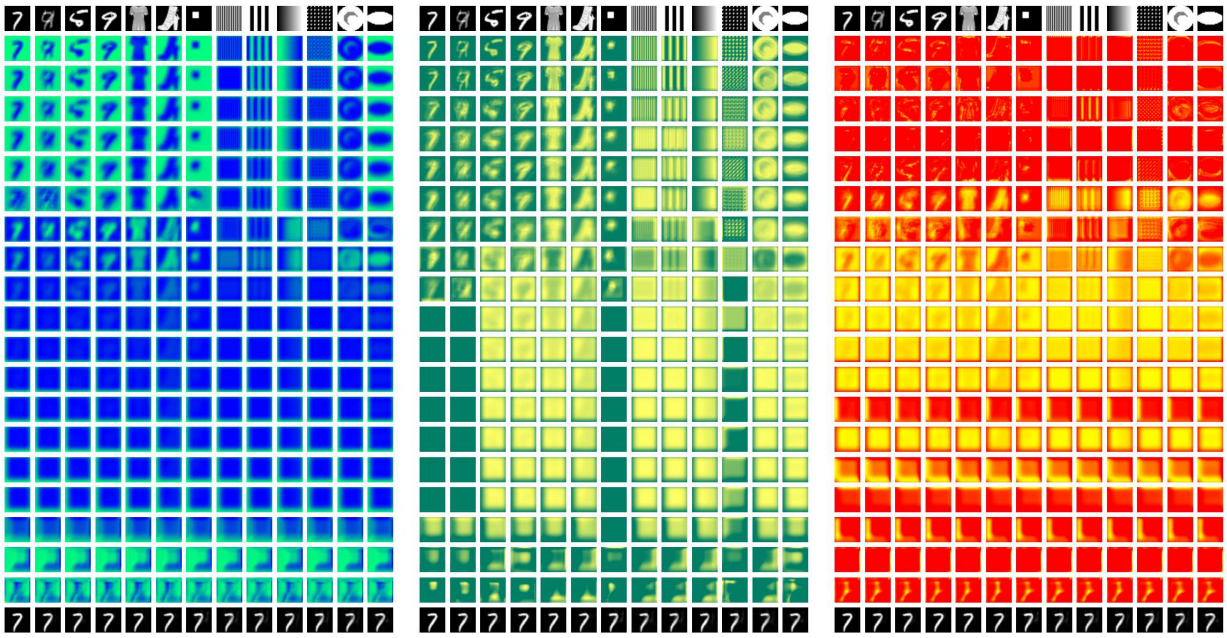

Figure 21: **Visualizing the intermediate layers of a trained 20-layer ConvNet.** The three subfigures show for each layer: 1) the top singular vector across the channels; 2) the channel that maximally correlate with the input image; 2) a random channel, respectively.

biases of a 5-layer ConvNet when the convolutional filter size changes. The visualization shows that the predictions become more and more blurry as the filter sizes grow. The heatmaps, especially the correlation to the identity function, are not as helpful in this case as the correlation metric is not very good at distinguishing images with different levels of blurry. As also discussed before, with extremely large filter sizes that cover the whole inputs, the ConvNets start to bias towards the constant function. Note our training inputs are of size 28×28, so 29×29 filter size allows all the neurons to see no less than half of the spatial domain from the previous layer. 57×57 receptive fields centered at any location within the image will be able to see the whole previous layer. On the other hand, the repeated application of *the same* convolution filter through out the spatial domain is still used (with very large boundary paddings in the inputs). So the ConvNets are *not* trivially doing the same computation as FCNs.

Figure 27 shows the correlation to the constant and the identity function when different numbers of convolution channels are used. The heatmap is consistent with the visualizations from Figure 15, showing that the 5-layer ConvNet fails to approximate the identity function when only three channels are used in each convolution layer. Furthermore, Figure 28 visualize the predictions of trained 3-channel ConvNets with various depths. The 3-channel ConvNets beyond 8 layers fail to converge during training. The 5-layer and the 7-layer ConvNets implement functions biased towards edge-detecting or countour-finding. But the 6-layer and the 8-layer ConvNets demonstrate very different biases. The potential reason is that with only a few channels, the random initialization does not have enough randomness to smooth out "unlucky" bad cases. Therefore, the networks have higher chance to converge to various corner cases. Figure 29 and Figure 30 compare the random initialization with the converged network for a 3-channel ConvNet and a 128-channel ConvNet. From the visualizations of the intermediate layers, the 128-channel ConvNet already behave more smoothly than the 3-channel ConvNet at initialization.

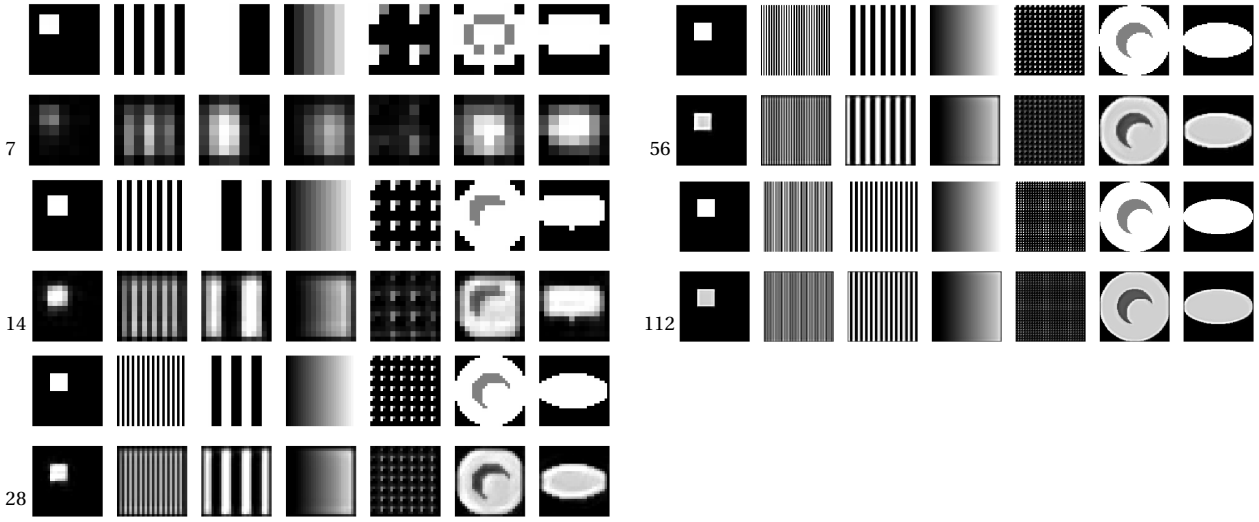

Figure 22: **Visualization of a 5-layer ConvNet on test images of different sizes.** Every two rows show the inputs and model predictions. The numbers on the left indicate the input image size (both width and height).

I Correlation vs MSE

Figure 31, Figure 32 and Figure 33 can be compared to their corresponding figures in the main text. The figures here are plotted with the MSE metric between the prediction and the groundtruth, while the figures in the main text uses the correlation metric. Each corresponding pair of plots are overall consistent. But the correlation plots show the patterns more clearly and has a fixed value range of [0, 1] that is easier to interpret.

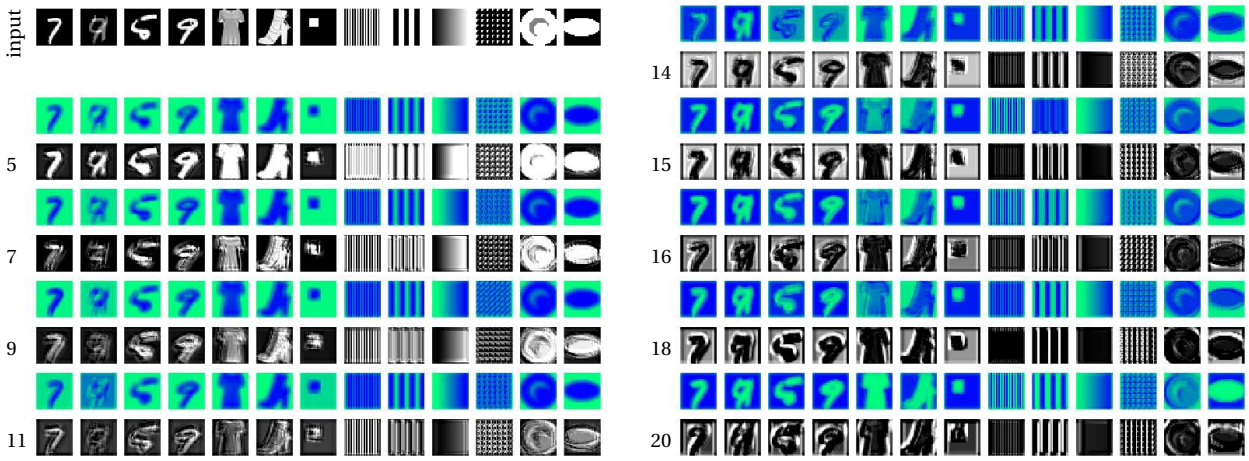

Figure 23: **Visualizing only the top two layers in trained networks.** The first row are the input images, which are directly fed into the top two layer of trained networks (skipping the bottom layers). The remaining rows shows the predictions from the top two layers of ConvNets, with the numbers on the left indicating their (original) depth. More specifically, each of the two top layers occupies one row. The colorful rows are the visualizations (as the top singular vector across channels) of the outputs of the second to the last layer from each network. The grayscale rows are the outputs of the final layer from each network.

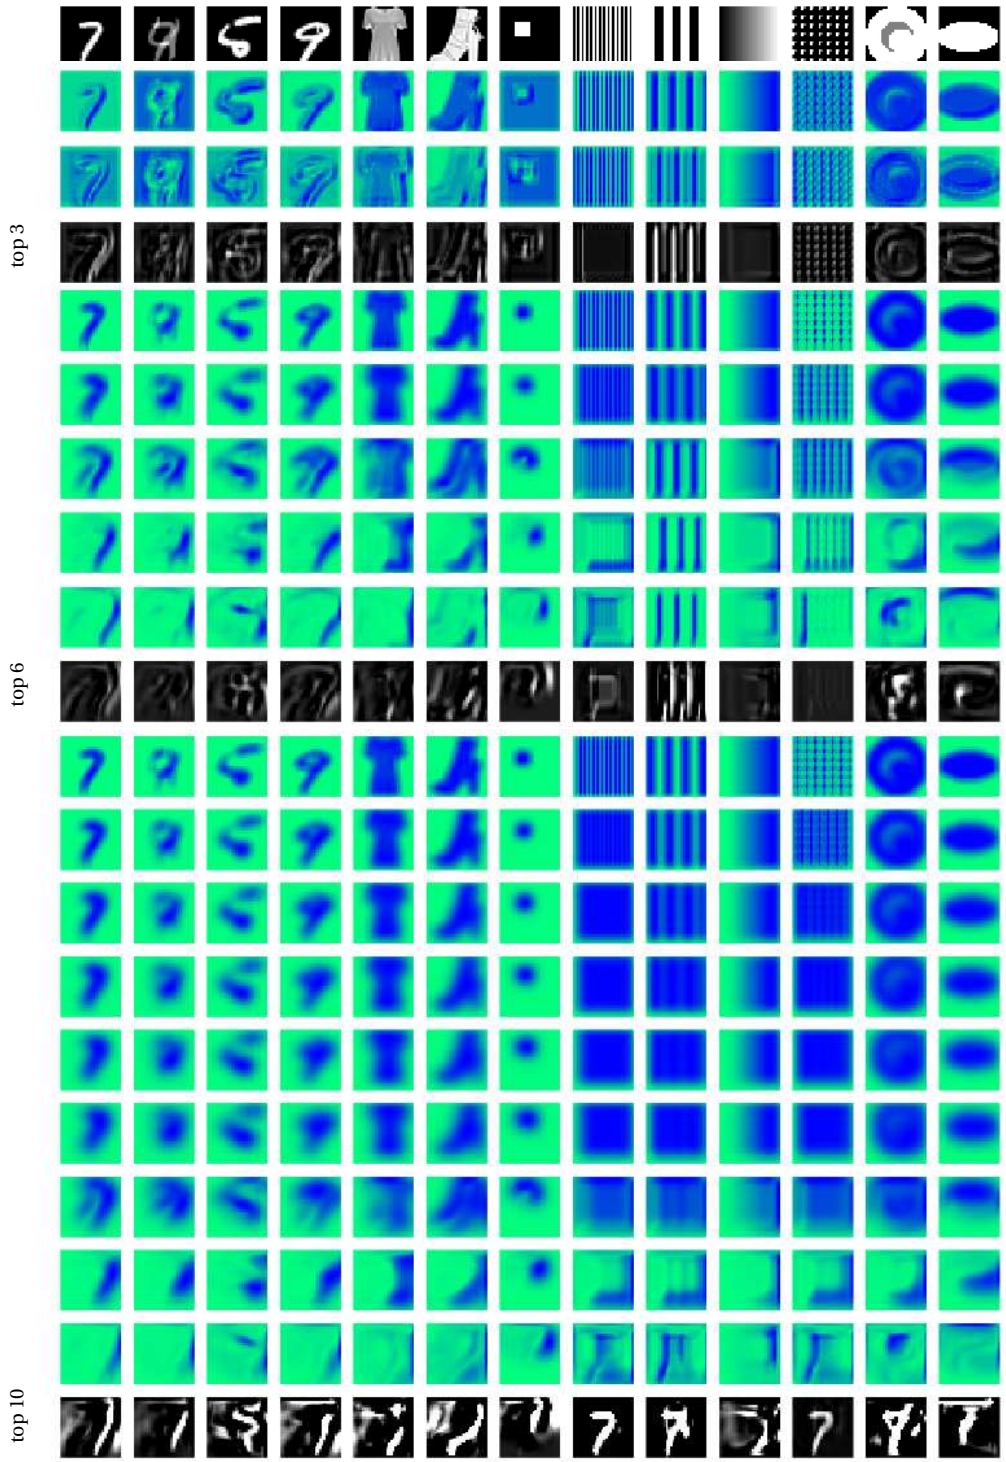

Figure 24: **Visualzing the top 3 layers, 6 layers and 10 layers of a 20-layer ConvNet.** Visualizations formatted in the same way as Figure 23.

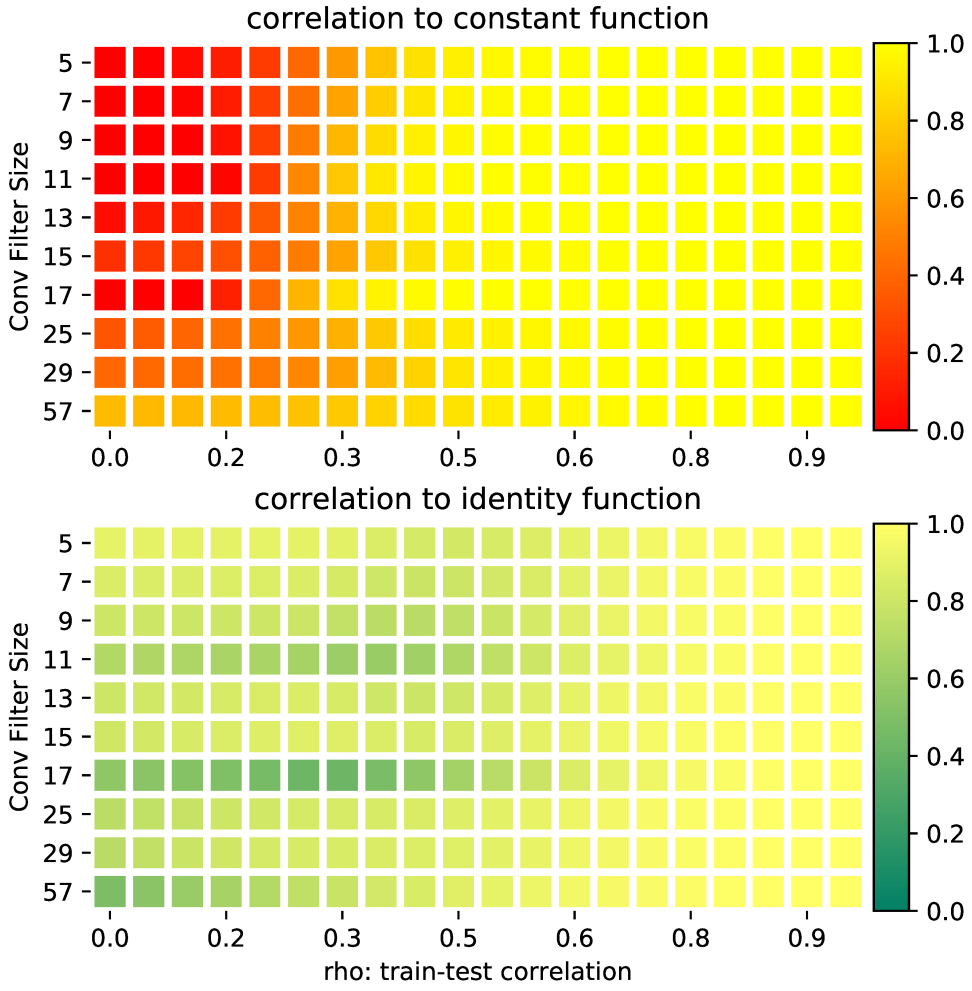

Figure 25: **Inductive bias of a 5-layer ConvNet with varying convolutional filter size.** The heatmap is arranged in the same way as Figure 7, except that the rows correspond to ConvNets filter sizes.

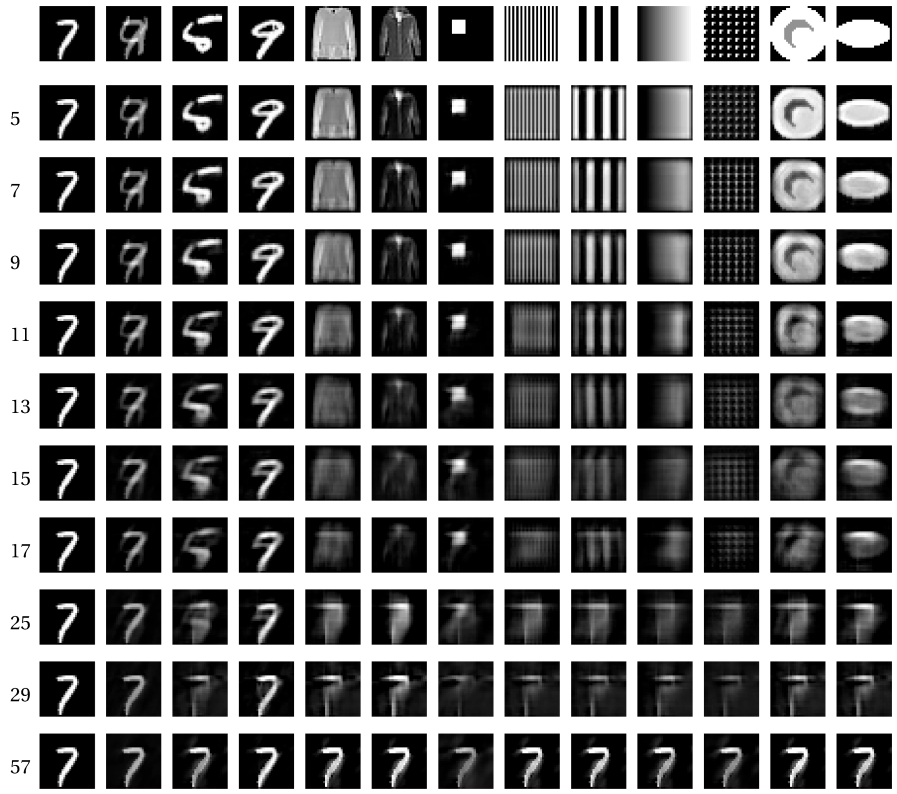

Figure 26: **Visualizing the predictions from ConvNets with various filter sizes.** The first row is the inputs, and the remaining rows are predictions, with the numbers on the left indicating the corresponding filter sizes. A shorter version of this is presented in Figure 14.

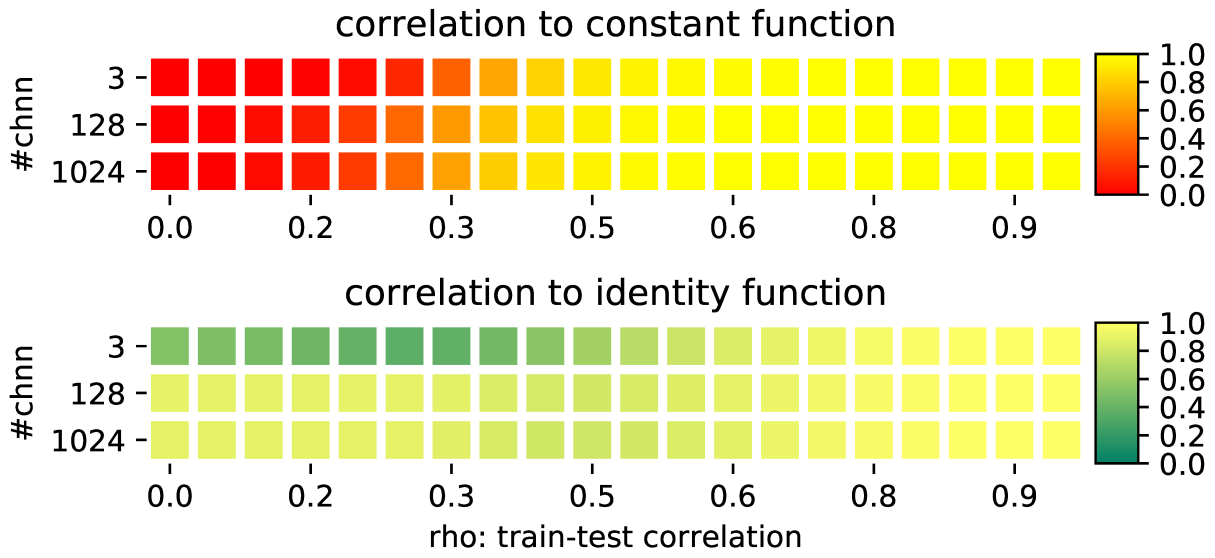

Figure 27: **Correlation to the constant and the identity function for different convolution channels in a 5-layer ConvNet.**

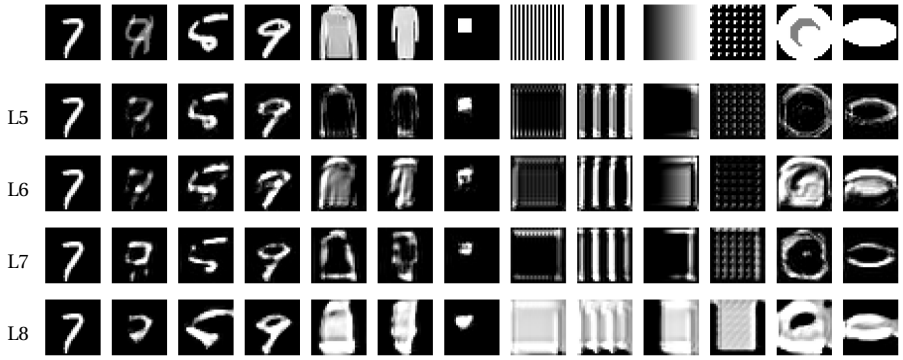

Figure 28: **Visualization predictions from ConvNets with 3 convolution channels and with various number of layers (numbers on the left).** The first row is the inputs, and the remaining rows illustrate the network predictions.

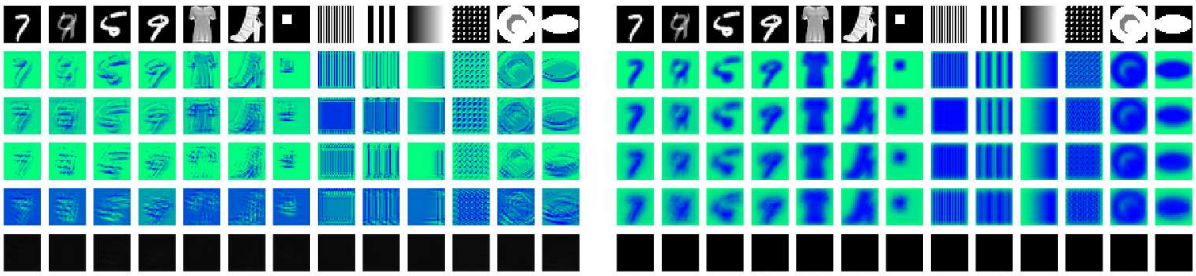

(a) 3 channels, random init

(b) 128 channels, random init

Figure 29: **Visualizing the randomly initialized models to compare two 5-layer ConvNets with 3 convolution channels per layer and 128 convolution channels per layer, respectively.** The subfigures visualize the predictions of intermediate layers of the two network at random initialization. The multi-channel intermediate layers are visualized as the top singular vectors.

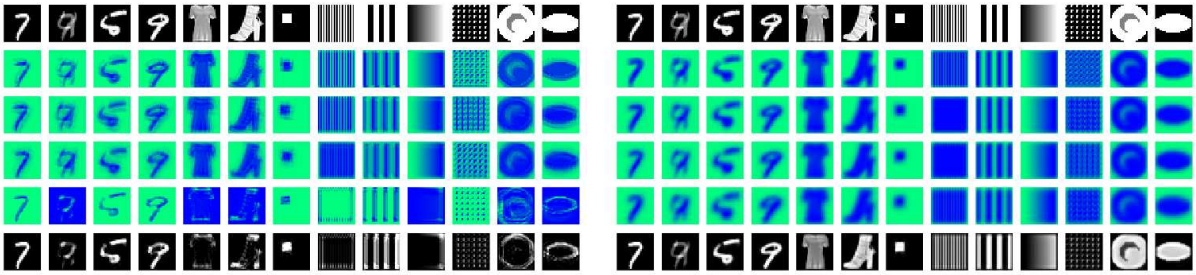

(a) 3 channels, after training

(b) 128 channels, after training

Figure 30: **Comparing two 5-layer ConvNets with 3 convolution channels per layer and 128 convolution channels per layer, respectively.** Layout is similar to Figure 29.

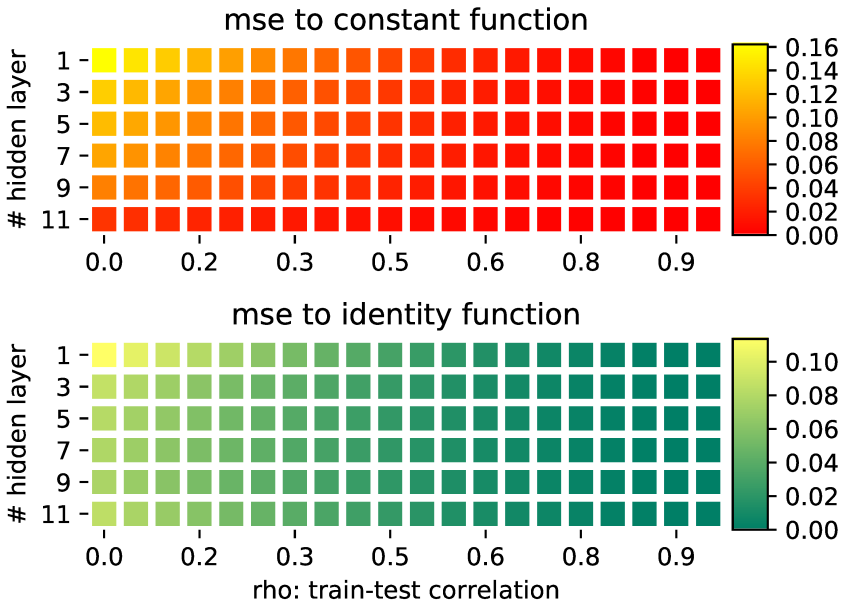

Figure 31: **Quantitative evaluation of linear FCNs.** The same as Figure 5(a), except MSE is plotted here instead of correlation.

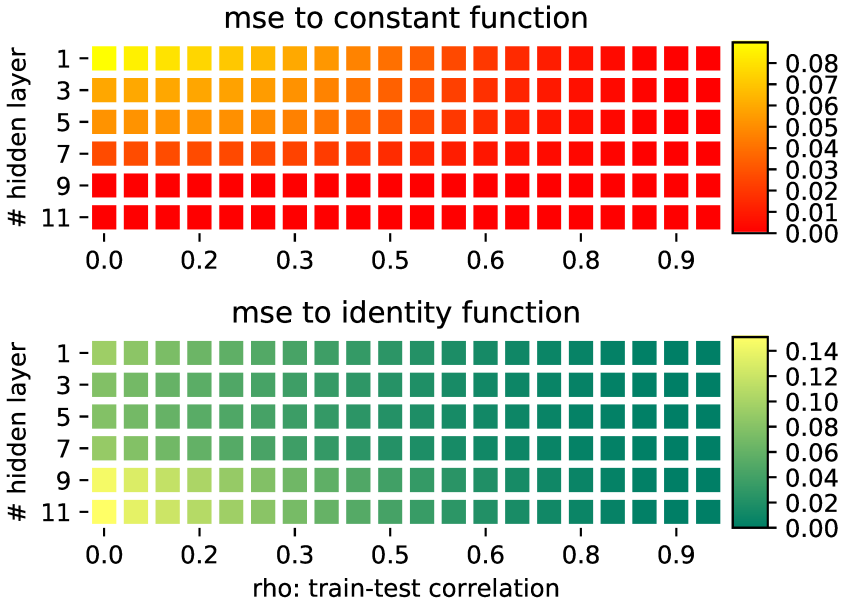

Figure 32: **Quantitative evaluation of ReLU FCNs.** The same as Figure 5(b), except MSE is plotted here instead of correlation.

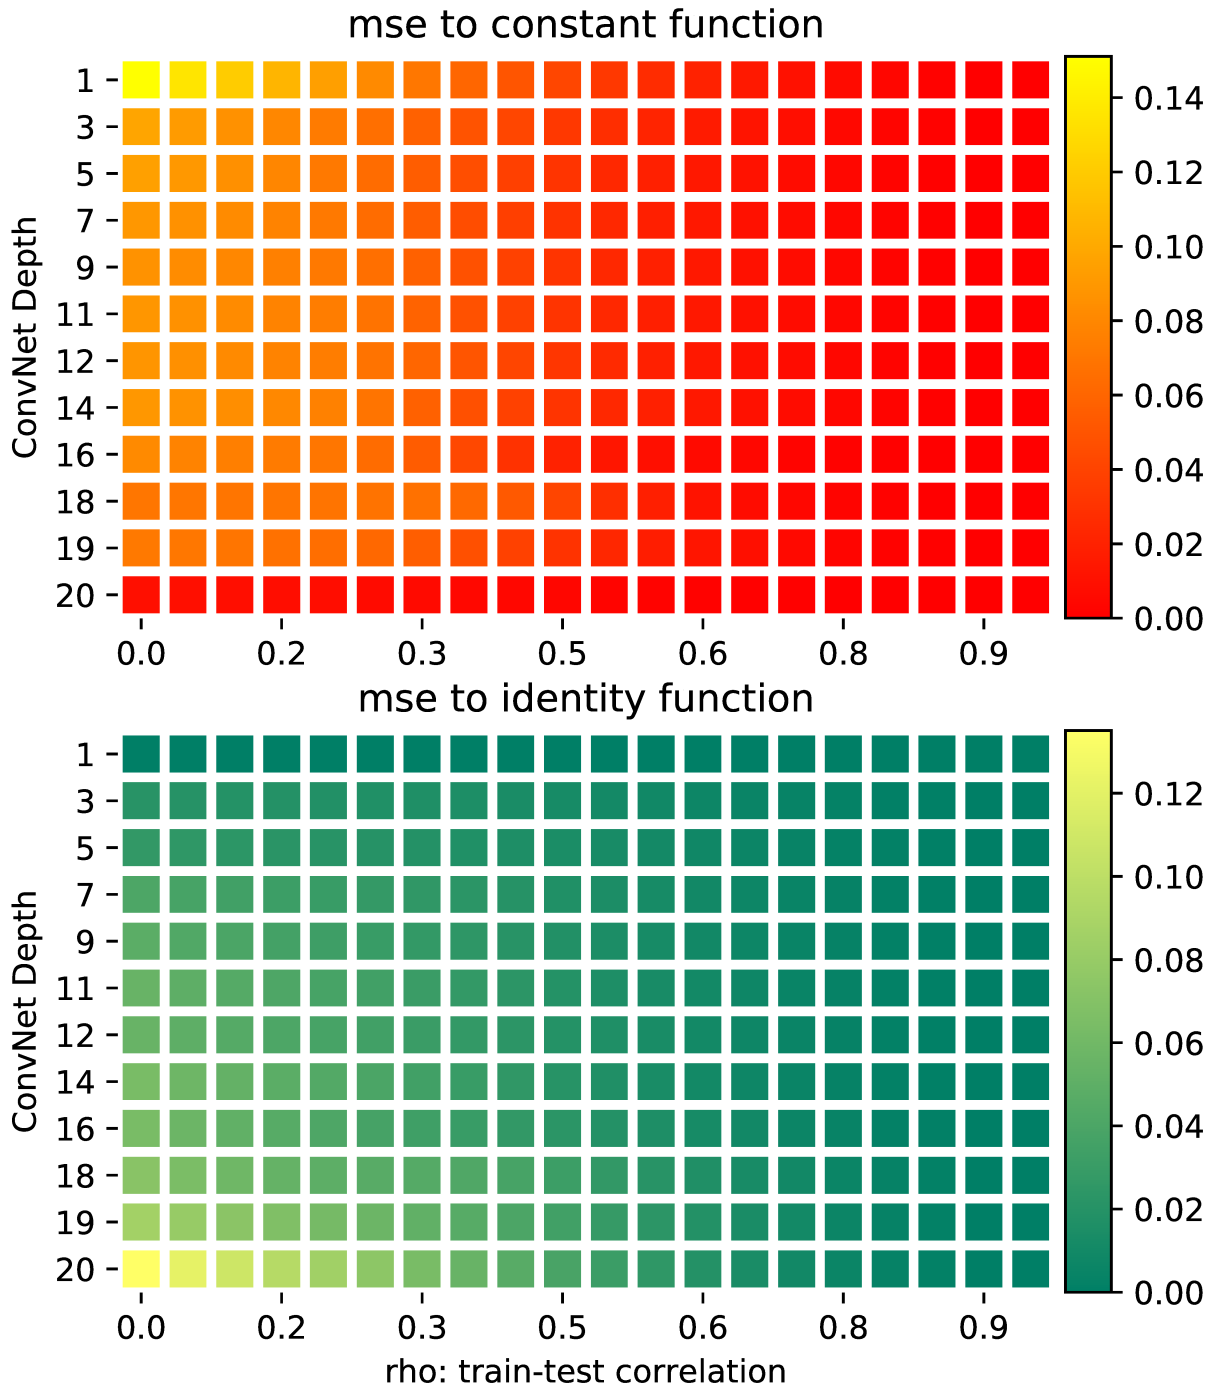

Figure 33: **Quantitative evaluation of ConvNets.** The same as Figure 7, except MSE is plotted here instead of correlation.