

Neural Representations Reveal Distinct Modes of Class Fitting in Residual Convolutional Networks

Michał Jamroz, Marcin Kurdziel

AGH University of Science and Technology, Krakow, Poland
mijamroz@agh.edu.pl, kurdziel@agh.edu.pl

Abstract

We leverage probabilistic models of neural representations to investigate how residual networks fit classes. To this end, we estimate class-conditional density models for representations learned by deep ResNets. We then use these models to characterize distributions of representations across learned classes. Surprisingly, we find that classes in the investigated models are not fitted in a uniform way. On the contrary: we uncover two groups of classes that are fitted with markedly different distributions of representations. These distinct modes of class-fitting are evident only in the deeper layers of the investigated models, indicating that they are not related to low-level image features. We show that the uncovered structure in neural representations correlate with memorization of training examples and adversarial robustness. Finally, we compare class-conditional distributions of neural representations between memorized and typical examples. This allows us to uncover where in the network structure class labels arise for memorized and standard inputs.

Introduction

Neural networks are vastly over-parametrized models. This flexible parameter space makes it quite difficult to answer basic questions about internal representations learned by these models. For a long time it was not even clear whether two identical networks trained on the same task learn similar sets of representations. That said, the picture begin to change in recent years with the introduction of new tools to analyse internal representations in neural networks. First, Raghu et al. (2017) and Morcos, Raghu, and Bengio (2018) proposed canonical correlation analysis-based similarity measures for neural representations. Subsequently, Kornblith et al. (2019) proposed a kernel-based similarity score and found significant representational similarity between networks trained to solve closely related tasks. Machinery for analysis of neural representations was then extended by Jamroz, Kurdziel, and Opala (2020), who proposed a non-parametric density model for features learned by neural networks. They used this model to show that memorizing networks learn more complex representations than networks that can exploit patterns in data.

So far neural representations were studied from the perspective of features learned by network layers. In this work

we extend this line of research to investigate *class representations* learned by neural networks. To this end, we estimate class-conditional distributions of neural representations and use them as proxies to the representations of classes learned by the network. This allows us to make several important contributions. First, we find that ResNets do not fit classes in a uniform way. On the contrary, we observe two distinct modes of class fitting evident in the deeper layers of these networks. We characterize class-conditional distributions of neural representations in the two groups of classes and propose a likely explanation for the observed differences. Next, we demonstrate that the uncovered structure in class representations translates to observable differences in memorization of input examples. To this end, we leverage tractable memorization measures recently proposed by Feldman and Zhang (2020). We also demonstrate that the two groups of classes differ in adversarial robustness. Finally, we leverage class-conditional distributions of neural representations to uncover where in the network structure classes are fit and to compare this process for memorized and typical inputs.

See Jamroz and Kurdziel (2022)¹ for the extended version of this paper, including the appendix and initial results for Vision Transformer and MLP-Mixer models. We provide the source code and our density and memorization estimates.²

Probabilistic Model for Class Representations

Our goal in this work is to characterize representations of classes in neural networks. While the class membership information is typically available for a large set of observations, namely the training set, combining it into representations of classes is a non-trivial task. The main hurdle here comes from the stochastic nature of neural representations. Specifically, a neural representation inferred for some input x can be seen as an outcome of sampling x from the data distribution. Clearly, a reasonable notion of a *class representation* should capture the outcome of this sampling. We therefore propose to use class-conditional distributions of inputs' representations as proxies to neural representations of classes. Concretely, we fit tractable density models to sets of neural representations and then use these models to characterize distributions of representations in classes.

¹<https://arxiv.org/abs/2212.00771>

²<https://github.com/mjamroz90/dnn-class-fitting>

We capture the distributions of neural representations with the hierarchical Bayesian model that was recently used by Jamroz, Kurdziel, and Opala (2020) to investigate networks that fit random labels. Jamroz, Kurdziel, and Opala used this model to characterize representations of kernels in convolutional layers. A neural representation in their work was therefore defined as a vector of kernel responses over a fixed sequence of inputs, averaged across spatial dimensions. Our goal, however, is not to investigate features learned by convolutional kernels, but to characterize—in distributional settings—representations of inputs from the learned classes. We therefore adopt a different construction for neural representations. Specifically, in modern variable-resolution convolutional models input to the classification head is frequently constructed by global average pooling of the output feature maps. We follow a similar construction for representations of inputs in our class-conditional density models. In particular, consider a network layer l with a sequence of convolutional kernels: $k_1^l, k_2^l, \dots, k_d^l$ —where d is the layer width—and let $\mathbf{x} \in C$ be an input observation from class C . We construct the representation of \mathbf{x} at layer l as the vector of respective kernel activations averaged across spatial dimensions:

$$nn_l(\mathbf{x}) = [\text{avg_pool}(k_1^l(\mathbf{x})), \text{avg_pool}(k_2^l(\mathbf{x})), \dots, \text{avg_pool}(k_d^l(\mathbf{x}))]. \quad (1)$$

Activation vectors for all inputs $\mathbf{x} \in C$ collectively form a set of observations that is then explained with the density model. Because all input observations are sampled from C , the model estimates the class-conditional distribution for C .

Density Model

The hierarchical Bayesian model used by Jamroz, Kurdziel, and Opala (2020) explains a set of observations $\mathbf{x} \in \mathcal{D}$ —e.g. a set of neural representations of some network inputs—with a mixture of multivariate normal distributions with an unknown number of components:

$$\begin{aligned} \alpha &\sim \text{Gamma}(1, 1), \\ G &| \alpha \sim \text{DP}(\text{NIW}(\boldsymbol{\theta}_0), \alpha), \\ \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k &\sim G, \\ \mathbf{x} &| \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \end{aligned} \quad (2)$$

The prior over component parameters is constructed as a Dirichlet Process (DP) with Normal-Inverse-Wishart (NIW) base distribution. A mixture of this form is consistent in total variation for a large family of continuous distributions (Ghosal and Van der Vaart 2017) and admits an efficient *collapsed Gibbs sampler* (CGS) for the posterior over component parameters (Neal 2000). Shortly, CGS constructs a Markov chain over assignments of observations to components: (c_1, c_2, c_3, \dots) . Under the model in Eq. (2), the conditional posterior predictive distribution $p(\mathbf{x}^* | \mathcal{D}, c_t)$ over a new observation \mathbf{x}^* given the set of assignments c_t has a closed-form solution (Jamroz, Kurdziel, and Opala 2020). This conditional predictive distribution can be leveraged to construct a tractable Monte Carlo approximation to the full posterior predictive distribution:

$$p(\mathbf{x}^* | \mathcal{D}) \approx \frac{1}{n} \sum_{t=1}^n p(\mathbf{x}^* | \mathcal{D}, c_t). \quad (3)$$

The average here is taken over the assignments sampled by CGS, i.e. over the Markov chain steps.

The closed-form solution for the predictive distribution conditioned on component assignments $p(\mathbf{x}^* | \mathcal{D}, c_t)$ gives a simple sampling strategy for the full posterior predictive: sample a component assignment c_t from the Markov chain and then sample \mathbf{x}^* from the induced conditional distribution $p(\mathbf{x}^* | \mathcal{D}, c_t)$. This sampling scheme can be used to compare the posterior predictive distribution with another distribution $q(\mathbf{x}^*)$ that has a tractable density. Specifically, the samples can be used to construct a Monte Carlo approximation to the relative entropy—or Kullback-Leibler (KL) divergence—from $q(\mathbf{x}^*)$ to $p(\mathbf{x}^* | \mathcal{D})$:

$$\begin{aligned} D_{\text{KL}}(p(\mathbf{x}^* | \mathcal{D}) || q(\mathbf{x}^*)) &= \\ &= \mathbb{E}_{\mathbf{x}^* \sim p(\mathbf{x}^* | \mathcal{D})} [\log p(\mathbf{x}^* | \mathcal{D}) - \log q(\mathbf{x}^*)] \\ &\approx \frac{1}{nm} \sum_{t=1}^n \sum_{s=1}^m [\log p(\mathbf{x}_{st}^* | \mathcal{D}, c_t) - \log q(\mathbf{x}_{st}^*)], \end{aligned} \quad (4)$$

where \mathbf{x}_{st}^* are sampled from $p(\mathbf{x}^* | \mathcal{D}, c_t)$. As proposed by Jamroz, Kurdziel, and Opala, the estimate in Eq. (4) can be used to quantify the complexity of neural representations. In this case the reference distribution $q(\mathbf{x}^*)$ is chosen to be the maximum entropy distribution that explains only the location and the scale of \mathcal{D} , i.e. the diagonal Gaussian with the mean and the variances estimated from \mathcal{D} .

In this work we use the divergence from the maximum entropy distribution to characterize the complexity of class-conditional distributions of neural representations. We also use Eq. (4) to compare posterior predictive distributions estimated for different classes. Specifically, let $p(\mathbf{x}^* | \mathcal{F})$ be the posterior predictive distribution under model in Eq. (2) induced by a set of neural representations \mathcal{F} . Because the density in this posterior predictive is tractable (Eq. (3)), we can use Eq. (4) to estimate the KL divergence from $p(\mathbf{x}^* | \mathcal{F})$ to $p(\mathbf{x}^* | \mathcal{D})$. To this end, we simply set the reference distribution $q(\mathbf{x}^*)$ to be the posterior predictive $p(\mathbf{x}^* | \mathcal{F})$.

Standard CGS exhibits relatively slow mixing. To speed-up the estimation of predictive distributions, in this work we sample posterior parameters using a variant of this algorithm with blocked Gibbs sampling (Jensen, Kjærulff, and Kong 1995). See Appendix for more details.

Memorization in Neural Networks

Recently, Feldman (2020) proposed a theoretical framework that explains the role of input memorization when fitting a long-tailed data distribution, i.e. a population of input examples characterized by many infrequent components. They consider a case where some of the examples in an n -element train set come from sub-populations with frequency below $1/n$. Such examples are statistically indistinguishable from noise or mislabelled data. Consequently, their fitting requires memorization. Importantly, Feldman showed that memorization in these settings is in fact necessary to minimize the generalization error. They also proposed a measure for the degree of input memorization. Specifically, given a learning algorithm \mathcal{A} and an example (\mathbf{x}, t) , Feldman defines the degree of memorization of (\mathbf{x}, t) as the probability of predicting t

for input \mathbf{x} when (\mathbf{x}, t) is observed during training, relative to the probability of the correct prediction when (\mathbf{x}, t) is absent from the training data. A score of 1.0 therefore corresponds to an example that is memorized in every training run, whereas a score of 0.0 corresponds to an example with no memorization. This definition assumes that \mathcal{A} is a stochastic learning algorithm, for example stochastic gradient descent. In a follow-up work Feldman and Zhang (2020) proposed a computationally-efficient approximation to this memorization score and used it to estimate memorization in two popular image classification benchmarks, namely CIFAR-100 and ImageNet. They found that in both cases convolutional neural networks memorize a non-trivial part of the train set.

Feldman and Zhang work opens an avenue for investigating mechanisms behind memorization in deep networks. In particular, previous research focused on models where memorization of specific examples was induced artificially, e.g. by corrupting the labels or the network input (Zhang et al. 2017; Arpit et al. 2017). Memorization estimates can, however, be used to investigate memorization in neural networks that learn from uncorrupted data. We use this opportunity to uncover the relationship between representations of classes in neural networks and the degree of input memorization. We also leverage class-conditional density models of neural representations to investigate where in the network structure classes arise for typical and memorized examples.

Experimental Setup

We conduct our experiments on two popular image datasets: CIFAR-100 (Krizhevsky 2009) and Mini-ImageNet (Vinyals et al. 2016). For CIFAR-100 we use the memorization scores published by Feldman and Zhang (2020).³ To avoid potential discrepancies between these estimates and our experimental conditions, we replicate important aspects of their training setup. Concretely, we investigate representations learned by a ResNet50 network that was trained by closely following hyper-parameters reported in (Feldman and Zhang 2020). To confirm our findings, we also repeat these experiments with another residual convolutional network, namely ResNet18. In this case we estimated the necessary memorization scores by following the Algorithm 1 in (Feldman and Zhang 2020). To this end, we trained 2000 network instances (*trials* in Algorithm 1 therein) in order to estimate the required probabilities of correct prediction. For the Mini-ImageNet experiments we collected images from the provided training, validation and test subsets, and then split them randomly into 50,000 training and 10,000 test examples. We then estimated memorization scores for the training images. These experiments were also carried out on the ResNet50 and ResNet18 architectures. However, to lessen the computational burden associated with estimation of memorization, we modified the original architectures by introducing a 3-pixel stride in the first convolutional layer. All other experimental details and training hyper-parameters match those used with the CIFAR-100 dataset. We provide memorization scores estimated for the CIFAR-100 and Mini-ImageNet datasets alongside this paper. Finally, it is worth noting that Feldman and Zhang

published memorization scores also for the ImageNet dataset. However, due to the cost of estimating class-conditional density models and related quantities on the ImageNet scale, we ultimately decided to focus on the Mini-ImageNet dataset.

ResNet architecture has four *stages* that correspond to progressively smaller spatial dimensions. We extract neural activations immediately after the summation operation at the end of each stage (See Appendix for details). In each case we collect activations for the entire train set and use them to calculate neural representations (Eq. 1). This gives us four sets of neural representations spaced evenly across the network depth. Importantly, our main findings come from the analysis of the representations in the last stage, i.e. the input to the classification head. We use representations from the intermediate stages only to show how class representations form across the network depth. While in principle we could analyse representations from each residual block, this would have no bearing on our main findings and would significantly increase the computational cost of experiments.

To obtain representations of classes we estimate the mixture model in Eq. (2) independently for each combination of the class label and the network stage. One difficulty in modelling ResNet representations comes from the width of this network: neural representations have up to 2048 dimensions. Estimating Gaussian mixture models in this many dimensions is expensive. To avoid this issue, we reduce the dimensionality of the collected representations via Singular Value Decomposition (SVD). In each case we retain enough dimensions to preserve most of the variance in the data set (See Appendix for details). It is worth noting that SVD preprocessing was used before in investigations of neural representations, e.g. by Raghu et al. (2017) and by Jamroz, Kurdziel, and Opala (2020). Importantly, Raghu et al. found that intrinsic dimensionality of neural representations—in their case the number of SVD directions needed to match the performance of a complete network—is much smaller than the number of neurons in the corresponding network layer, which justify the SVD pre-processing step.

To estimate the posterior distributions in the density models (Eq. (2)) we perform, in each case, 400 block collapsed Gibbs sampler steps (i.e. passes over the dataset). We use a block size of $b = 4$ observations for both datasets. Afterwards, we discard initial 320 Gibbs steps and include every 4th of the remaining steps in the Monte Carlo estimates (Eqs. (3) and (4)). Following Jamroz, Kurdziel, and Opala (2020), we put a data-derived, weakly-informative prior on the parameters of Gaussian components. To this end, we adopt the values of prior hyper-parameters (θ_0 in Eq. (2)) used in that work.

How Residual Networks Fit Classes

We begin our analysis by characterizing representations of classes learned from the two image datasets. First, for each class C and each network stage we calculate mean and standard deviation of class-conditional log-densities (Eq. 3) of examples $\mathbf{x} \in C$. Results for the third and fourth stage are reported in Fig. 1. Results for the first two stages are in the Appendix. Already, we see an unexpected structure: class representations in the last two stages display an almost phase

³Available at: <https://pluskid.github.io/influence-memorization>

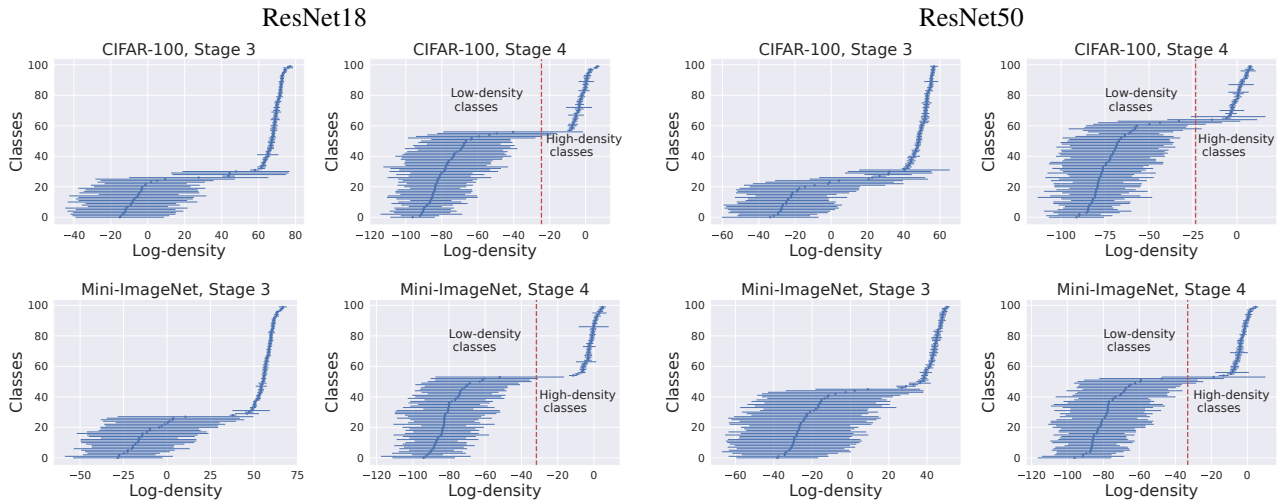


Figure 1: Mean and standard deviation of class-conditional log-densities estimated for inputs from each class. Classes on the vertical axis are sorted according to the mean class-conditional density.

transition-like change in class-conditional log-densities. That is, we observe two groups of classes that cluster around vastly different density values. Equally striking, classes in these groups typically differ by an order of magnitude in the standard deviation of the estimated class-conditional log-densities. Specifically, estimates for examples from the classes clustered around the higher class-conditional density typically have an order of magnitude lower variance. These observations agree between the two image datasets and the two architectures used in the experiment. Importantly, we did not find a similar structure in the representations from the first two stages, indicating that the observed phenomenon is not related to low-level features of the inputs (See Appendix). We also replicated this experiment using a plain (i.e. non-residual) convolutional network. This time we did not observe distinct modes of class fitting in any of the network layers (See Appendix). This implies that the observed structure is not simply a product of the datasets used in the experiments, but is related to the residual architecture. In summary, we observe two distinct modes of class fitting in the last two stages of ResNet models.

We now focus on the highest-level class representations, i.e. representations constructed by the last stage of the network. First, we split classes into two groups according to the estimated mean log-density values, namely *low-* and *high-density* classes (marked on Figure 1, stage 4). To this end, we sort classes according to their mean log-densities and find interval (C_i, C_{i+1}) in this sequence with the largest mean log-density difference. We use center of this interval to assign classes to the two groups. Note that we group entire classes, not individual input examples.

The most straightforward explanation for the observed differences between the uncovered groups could be that the high-density classes simply have more spatially compact distributions of neural representations. In other words, examples from these classes could have more similar high-level neural representations. Our results show that this simple explana-

tion is incorrect. Specifically, for each class C we calculated the mean distance $\|\text{svd}(nn_l(\mathbf{x})) - \text{svd}(nn_l(\mathbf{y}))\|$ between neural representations of examples $\mathbf{x}, \mathbf{y} \in C$ (Eq. (1), after dimensionality reduction with SVD). We also estimated the complexity of its class representation, namely the relative entropy of its posterior predictive distribution $p(\mathbf{x}^* | \mathbf{x}^* \in C)$ from the reference maximum entropy distribution $q(\mathbf{x}^*)$ (Eq. (4), and the paragraph below). Results are reported in Fig. 2. Clearly, the high-density classes are *not* more spatially compact at the neural representation level than the low-density classes. However, their posterior predictive distributions have vastly larger complexity than the posterior predictive distributions of the low-density classes. In other words, representations of the high-density classes are vastly more non-Gaussian. Together with the high average class-conditional density of their examples, these results point to a different explanation: in representations of high-density classes, the probability mass is concentrated in a number of compact but spatially separated modes. This suggests that at the neural representation level high-density classes are formed by a collection of compact components.⁴

Class Representations Correlate with Memorization and Adversarial Robustness

Our results so far uncover an unexpected structure in the representations of learned classes. An immediate question that follows from this observation is: does this structure correlate with some phenomena observed in neural networks? We identify two such phenomena: memorization of input examples and adversarial robustness.

⁴One could ask in this place: is this larger number of components observed in the posterior over parameters of model (2)? We do observe larger number of components sampled by CGS for high-density classes. However, we do not report these numbers, as—for technical reasons—estimation of component counts is fragile. See Appendix for details.

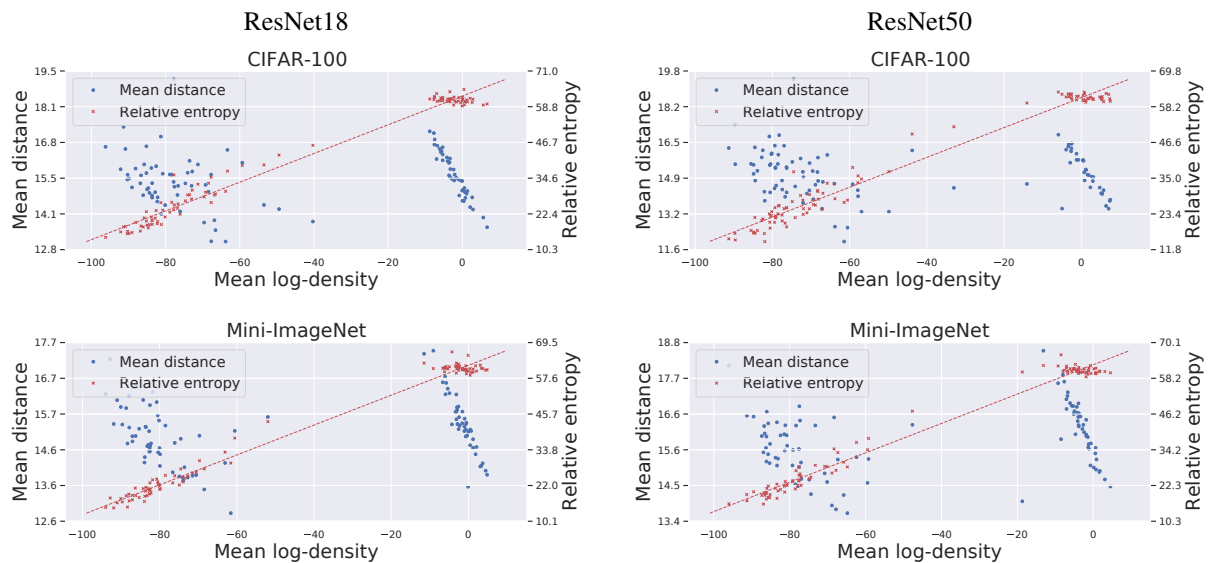


Figure 2: Relative entropies of class-conditional distributions of neural representations and mean intra-class distances between representations. Each class corresponds to one blue and one red mark. Dashed line: least-squares fit to relative entropies.

To demonstrate the relationship between input memorization and class representations, we sorted the training examples according to their estimated class-conditional log-density and then split this ordering into 50 equally sized bins. We then calculated for each bin the mean and the standard deviation of memorization scores and class-conditional log-densities of examples assigned to it. We also calculated the ratio of examples belonging to the low- and the high-density classes. Finally, for each bin we trained a separate ResNet model without using the training examples assigned to that bin. We then calculated the F-score of this model on the examples from the bin. Results are reported in Fig. 3.

Even though memorization scores exhibit large variance, averaging across density bins uncovers a difference in their distribution between the low- and the high-density classes. More precisely, up to the transition point between the low- and the high-density classes, the degree of memorization of an input examples decreases—on average—with increasing class-conditional log-density. This agrees with an intuition that increasing class-conditional density translates to an input example that is increasingly similar at the neural representation level to many other class members. However, the trend changes abruptly with the transition to the high-density classes. The transition is marked by an abrupt increase in average memorization of input examples, which subsequently remains largely independent of the class-conditional density. These observations are corroborated with the calculated F-scores, which are in close agreement with mean memorization scores.

At first glance, results in Fig. 3 may seem counter-intuitive. However, we argue that they corroborate the long-tail hypothesis put forward by Feldman (2020). Concretely, our results indicate that distributions of neural representations in the high-density classes have many distinct components. That is,

the high-density classes are formed from smaller subpopulations of examples with distinct neural representations. The frequency of each such subpopulation in the training data will be below the class frequency. Feldman’s hypothesis suggests that in order to minimize the generalization error, the learning model may need to memorize some of the examples from these subpopulations. And indeed, we observe an increase in the degree of memorization when the distribution of input examples switches from the low- to the high-density classes.

Compact and spatially separated components in the high-density classes should—intuitively—be less robust to an adversarial attack. In particular, a relatively small input perturbation may move the representation of the attacked example outside of its component. This could be verified by comparing low- and high-density classes w.r.t the robustness against a selection of adversarial attacks. However, given the vast number of attacks proposed so far, we opt to explore this hypothesis in an attack-agnostic way. To this end, we evaluate the performance of a provable adversarial defense in function of the class-conditional density. Concretely, we take the training examples split into density bins and for each bin train a smoothed classifier proposed by Cohen, Rosenfeld, and Kolter (2019) without using the examples from the bin. We use ResNet networks as the base (smoothed) models. We then evaluate the CERTIFY procedure (Cohen, Rosenfeld, and Kolter 2019) for the examples from the held-out bin (See Appendix for details). The procedure either abstains from prediction or returns a certified radius r and the predicted class. Importantly, Cohen, Rosenfeld, and Kolter proved that if not abstaining, CERTIFY returns with high-probability the class that will be predicted by the smoothed classifier for inputs no further than r from the certified example. Note, however, that the smoothed classifier need not agree with the base classifier. In this sense it trades classification accuracy

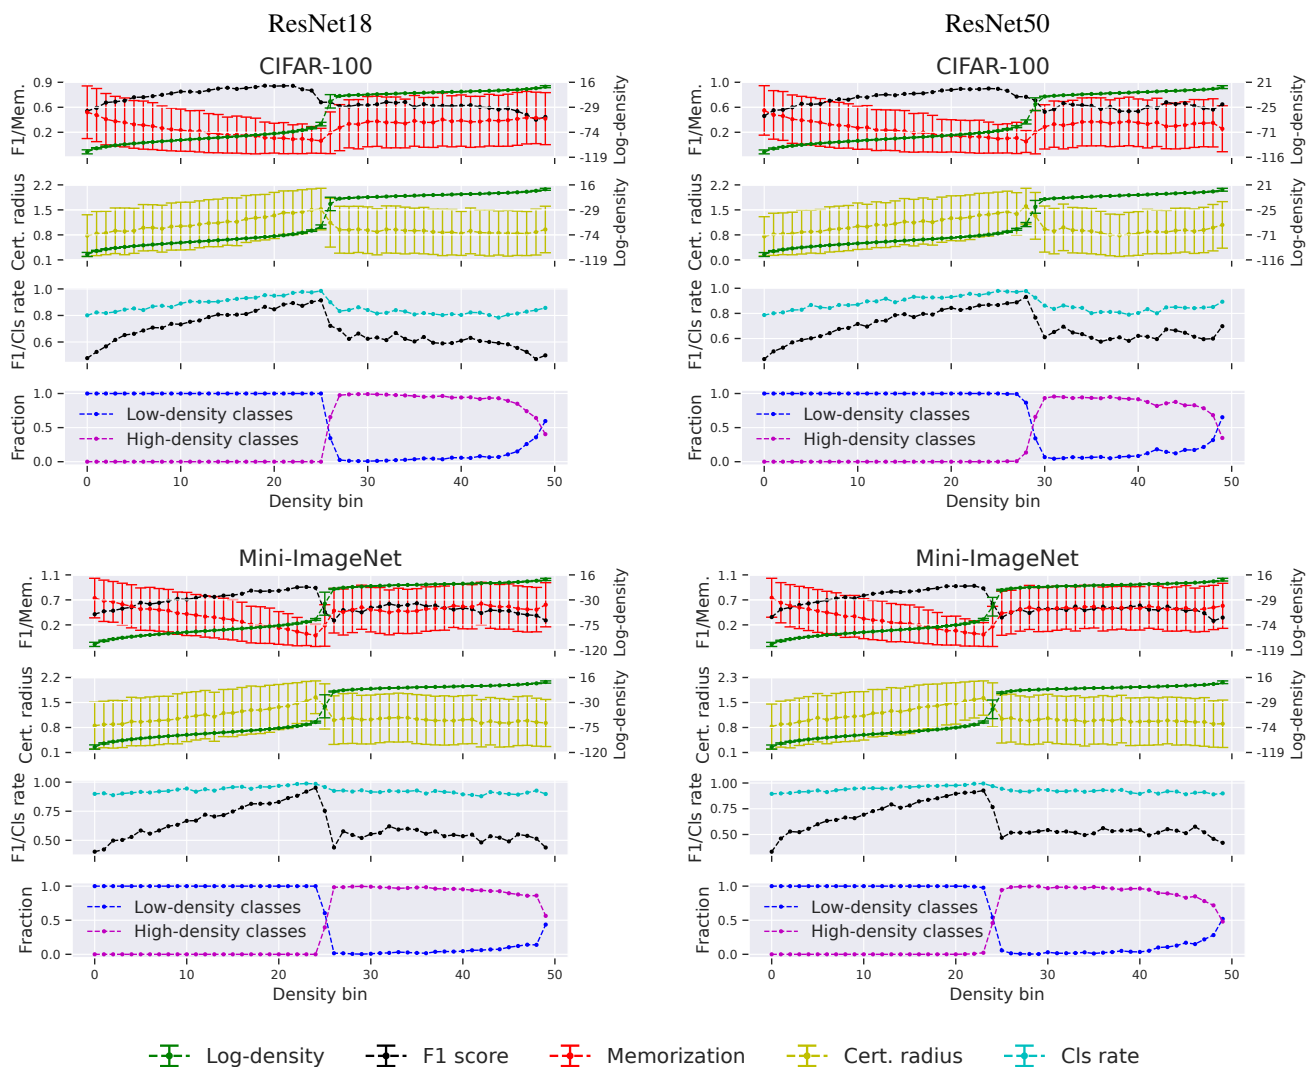


Figure 3: Relationship between class conditional log-density, input memorization and adversarial robustness. Low/high-density fraction: fraction of points in a bin from the low-/high-density classes.

for provable robustness.

In Fig. 3 we report per-bin certification radius (mean and standard deviation), fraction of inputs for which CERTIFY did not abstain (classification rate) and F-score of the smoothed classifier. Results confirm that high-density classes correlate with lower adversarial robustness: transition from the low- to the high-density regime coincide with abrupt decrease in certified radii and F-scores of the smoothed classifier. CERTIFY is also slightly more likely to abstain in high-density classes.

Where Neural Networks Fit Classes

We also use class-conditional distributions of neural representations to uncover where in the network architecture classes are fit and to show to what extent memorization influences this process. To this end, we selected from each dataset examples with memorization score above 0.9. Additionally, we selected approx. equal number of examples with the lowest memorization estimates and an

equal number of randomly chosen examples. We estimated the class-conditional posterior predictive distributions separately for the three chosen subsets of examples. Next, we estimated the KL divergences between predictive densities of every class pair (Eq. (4)). Note that for a pair of classes (P, Q) we estimate the KL divergence from $p(nn_l(\mathbf{x}) | \mathbf{x} \in P)$ to $p(nn_l(\mathbf{x}) | \mathbf{x} \in Q)$ and the divergence from $p(nn_l(\mathbf{x}) | \mathbf{x} \in Q)$ to $p(nn_l(\mathbf{x}) | \mathbf{x} \in P)$. The asymmetry of the KL divergence therefore does not alter the conclusions of this experiment. To recover meaningful predictive densities, we restrict this analysis to classes that have at least 100 representatives in all three subsets.

Figure 4 reports mean estimated between-class KL divergences, together with one standard deviation intervals. Additional results are reported in the Appendix. Clearly, input memorization is not fixed to some specific parts of the network: class-fitting progresses similarly for memorized and typical examples, although memorized examples induce




Figure 4: Kullback-Leibler divergences between class-conditional predictive densities.

slightly less distinct class representations. Importantly, class representations are formed mostly between the second and the third stage of the ResNet model. This agrees with the observation that distinct groups of classes are evident in the third and the fourth network stage, but not in the first two stages.

Related Work

Several recent works investigate representations learned by neural networks. The first algorithm to uncover similarities between representations in these models was proposed by Raghu et al. (2017). In crux, they use SVD to remove nuisance factors from network activations, and then uncover shared representations via canonical correlation analysis. Subsequently, Morcos, Raghu, and Bengio (2018) refined this approach with a novel method for combining canonical directions, while Kornblith et al. (2019) proposed a Hilbert-Schmidt Independence Criterion-based metric for comparing neural representations. Next, Jamroz, Kurdziel, and Opala (2020) proposed probabilistic models for representations learned by kernels in convolutional networks. They showed that networks that memorize random labels learn significantly more complex representations than generalizing networks. Note that while we use a similar set of of probabilistic tools, our goal is not to characterize distributions of features learned by network units. Instead, we use class-conditional probabilistic models to capture a notion of class representations in neural networks.

A growing line of research touch on the memorization of input examples by neural networks. Zhang et al. (2017) demonstrate that neural networks can easily fit a dataset with randomly permuted labels, often despite explicit regularization during training. Arpit et al. (2017) compare networks that fit corrupted labels and networks trained on uncorrupted data with respect to the way they fit input examples. Zhang et al. (2018) proposes a new augmentation method that encourages the model to interpolate linearly between training examples. Their results suggest that such augmentation can prevent neural networks from memorizing random datasets and make them less prone to adversarial attacks. Effects of memorization were also covered by Neyshabur et al. (2017), who derived several metrics for the degree of memorization by bounding the Lipschitz constant in the model’s transformation. Finally, Feldman (2020, 2021); Feldman and Zhang (2020) propose an alternative perspective on memorization. They consider a learning task where data is assumed to follow

a long-tailed distribution with many infrequent components. Feldman shows that in this scenario minimization of the generalization error requires memorizing some of the examples.

Szegedy et al. (2014) demonstrated vulnerability of neural networks to an adversarial attack. Subsequently, Goodfellow, Shlens, and Szegedy (2015) proposed a simple method for constructing adversarial examples. These results prompted significant research efforts on adversarial attacks, defence strategies and certification of models’ predictions—see Huang et al. (2020) for a recent survey.

Conclusions

In this work we used Bayesian mixture models with unknown number of components to investigate representations of classes in residual convolutional networks. Our main finding is that classes in investigated neural models are fit in two distinct ways. Namely, we uncover a group of classes in which high-level neural representations appear to form compact and spatially separated components. We showed that examples from these classes are memorized to a higher degree than examples with intermediate class-conditional log-density estimates. We also showed that these classes are less robust to an adversarial attack. Finally, we used class-conditional density models to uncover where in the network structure class representations are formed for typical and memorized examples.

Our findings gives further experimental support for the perspective on memorization proposed in (Feldman 2020; Feldman and Zhang 2020). In particular, Feldman (2020) argues that memorization of examples from infrequent components is necessary to minimize generalization error in learning tasks where data comes from a long-tailed distribution. We do observe increased memorization in classes that appear to be formed from distinct—at the neural representation level—subpopulations of examples. While we do not claim to have a theoretical explanation for why representations of these classes are so strikingly different than representations of low-density classes, we point out that current convolutional architectures evolved largely through experimental optimization of performance on several image classification benchmarks. Feldman and Zhang (2020) identify cases in these benchmarks where memorization of certain inputs improve predictions on evaluation sets. Therefore, it may simply be that convolutional networks used today were unwittingly optimized to fit classes the way we observe in this work.

Acknowledgments

Research presented in this paper was supported by funds assigned to AGH University of Science and Technology by the Polish Ministry of Education and Science. This research was supported in part by PL-Grid Infrastructure.

References

- Arpit, D.; Jastrzebski, S.; Ballas, N.; Krueger, D.; Bengio, E.; Kanwal, M. S.; Maharaj, T.; Fischer, A.; Courville, A. C.; Bengio, Y.; and Lacoste-Julien, S. 2017. A Closer Look at Memorization in Deep Networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, 233–242.
- Cohen, J. M.; Rosenfeld, E.; and Kolter, J. Z. 2019. Certified Adversarial Robustness via Randomized Smoothing. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, 1310–1320. PMLR.
- Feldman, V. 2020. Does Learning Require Memorization? A Short Tale about a Long Tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, 954–959. Association for Computing Machinery.
- Feldman, V. 2021. Does Learning Require Memorization? A Short Tale about a Long Tail. *arXiv preprint arXiv:1906.05271*.
- Feldman, V.; and Zhang, C. 2020. What Neural Networks Memorize and Why: Discovering the Long Tail via Influence Estimation. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Ghosal, S.; and Van der Vaart, A. 2017. *Fundamentals of nonparametric Bayesian inference*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015*.
- Huang, X.; Kroening, D.; Ruan, W.; Sharp, J.; Sun, Y.; Thamo, E.; Wu, M.; and Yi, X. 2020. A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability. *Computer Science Review*, 37(100270).
- Jamroz, M.; Kurdziel, M.; and Opala, M. 2020. A Bayesian Nonparametrics View into Deep Representations. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 1440–1450.
- Jamroz, M.; and Kurdziel, M. 2022. Neural Representations Reveal Distinct Modes of Class Fitting in Residual Convolutional Networks. *arXiv preprint arXiv:2212.00771*.
- Jensen, C. S.; Kjærulff, U.; and Kong, A. 1995. Blocking Gibbs sampling in very large probabilistic expert systems. *Int. J. Hum. Comput. Stud.*, 42(6): 647–666.
- Kornblith, S.; Norouzi, M.; Lee, H.; and Hinton, G. E. 2019. Similarity of Neural Network Representations Revisited. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, 3519–3529.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Morcos, A. S.; Raghu, M.; and Bengio, S. 2018. Insights on representational similarity in neural networks with canonical correlation. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, 3-8 December 2018, Montréal, Canada*, 5732–5741.
- Neal, R. M. 2000. Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Stat.*, 9(2): 249–265.
- Neyshabur, B.; Bhojanapalli, S.; McAllester, D.; and Srebro, N. 2017. Exploring Generalization in Deep Learning. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 5947–5956.
- Raghu, M.; Gilmer, J.; Yosinski, J.; and Sohl-Dickstein, J. 2017. SVCCA: Singular Vector Canonical Correlation Analysis for Deep Learning Dynamics and Interpretability. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, 6076–6085.
- Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I. J.; and Fergus, R. 2014. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014*.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Kavukcuoglu, K.; and Wierstra, D. 2016. Matching Networks for One Shot Learning. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10 2016, Barcelona, Spain*, 3630–3638.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2017. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017*.
- Zhang, H.; Cissé, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2018. mixup: Beyond Empirical Risk Minimization. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018*.