Network-to-Network Regularization: Enforcing Occam's Razor to Improve Generalization

Anonymous Author(s)
Affiliation
Address
email

Abstract

What makes a classifier have the ability to generalize? There have been a lot of important attempts to address this question, but a clear answer is still elusive. Proponents of complexity theory find that the complexity of the classifier's function space is key to deciding generalization, whereas other recent work reveals that classifiers which extract invariant feature representations are likely to generalize better. Recent theoretical and empirical studies, however, have shown that even within a classifier's function space, there can be significant differences in the ability to generalize. Specifically, empirical studies have shown that among functions which have a good training data fit, functions with lower Kolmogorov complexity (KC) are likely to generalize better, while the opposite is true for functions of higher KC. Motivated by these findings, we propose, in this work, a novel measure of complexity called Kolmogorov Growth (KG), which we use to derive new generalization error bounds that only depend on the final choice of the classification function. Guided by the bounds, we propose a novel way of regularizing neural networks by constraining the network trajectory to remain in the low KG zone during training. Minimizing KG while learning is akin to applying the Occam's razor to neural networks. The proposed approach, called network-to-network regularization, leads to clear improvements in the generalization ability of classifiers. We verify this for three popular image datasets (MNIST, CIFAR-10, CIFAR-100) across varying training data sizes. Empirical studies find that conventional training of neural networks, unlike network-to-network regularization, leads to networks of high KG and lower test accuracies. Furthermore, we present the benefits of N2N regularization in the scenario where the training data labels are noisy. Using N2N regularization, we achieve state-of-the-art performance on MNIST with corrupted training labels and competitive performance on CIFAR-10. These findings illustrate the many benefits obtained from imposing a function complexity prior like Kolmogorov Growth during the training process.

1 Introduction and Motivation

On the surface, the problem of learning to generalize well over unseen data seems an impossible task. Classification is inherently a problem in function estimation, and the finite information that the training data samples impart seems hardly enough to be able to correctly guess the behaviour of the function over the unseen data samples outside the training set. However, assumption of a structured ground truth label function, which is the unknown function which generates the ground truth label for any datapoint, leads to more optimistic outlook on the problem. Without the assumption of structuredness in the ground truth, learning is not guaranteed, as was observed in the no free lunch theorem [1]. As shown in that work, there is no universal learning algorithm which can generalize well for possible choices of ground truth label functions.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

Since one does not have any control over the true nature of the ground truth function in any classification problem, the other important parameter that decides the ability of a classifier to generalize, is the complexity of its function space itself. Over the past few decades, there have been multiple attempts at bounding the generalization error on the basis of various complexity measures [2, 3, 4] of the classifier's function space $\mathcal{F}$. The overall results of these theoretical developments indicate that generalization is primarily governed by metrics that are proportional to the *size* of $\mathcal{F}$. Examples of complexity metrics in this regard include Rademacher complexity, VC-Dimension, and Local forms of Rademacher Complexity (for more tight bounds). However, as observed in [5] these metrics lead to relatively loose bounds on the generalization performance for deep neural networks. This is because deep neural networks have high complexity spaces with flexibility to learn any label assignment on any set of training datapoints. This was perhaps most clearly observed in [6], where even after random label assignments on the training data samples, the networks still were able to fit the training data labels with no errors.

These studies lead to a natural question, which is, other than the metrics which relate to the size of the function space of a classifier, what other factors contribute towards its ability to generalize? State-of-the-art deep neural networks, as observed above, will yield very high complexity spaces, which does not explain their remarkable ability to generalize well in complex high-dimensional supervised classification problems, such as in vision. We note that metrics such as Rademacher complexity or VC dimension, which relate to whole function spaces $\mathcal{F}$ (or subspaces within $\mathcal{F}$ which fit the training data well), essentially assign the same generalization gap to all functions within $\mathcal{F}$ (or the subspace within $\mathcal{F}$). However, there has been a longstanding understanding of the fact that usually among all functions which fit the training data well, simpler functions are expected to generalize better. This is an example of the Occam's Razor principle, which states that among all hypotheses that explain a phenomenon, the simplest hypothesis is preferred [7]. In the context of neural networks, even for very high complexity deep neural network function spaces, there will always be some weight configurations of a deep neural network, which yield *simpler* functions. For example, all network weight configurations yielding input-output functions which can still be efficiently approximated by smaller, shallower networks, could be considered to be *simpler*. An extreme example of this would be where we assign all weight values to zero within a deep neural network. The function that results from this weight configuration is essentially the constant function $f(X) = 0$. This observation points to the fact that even within a function space of a deep CNN, not all functions are *equally complex*, as some of them can still be approximated by shallower networks.

To that end, a primary objective of this paper is to probe complexity measures which enable us to assign a level of complexity to individual functions within a function space. Subsequently, our objective would be to steer the learning process towards network configurations which yield less complex input-output functions. An important work which explored similar directions is [8], where metrics from algorithmic complexity were developed to bias the learning process towards simpler functions. In this regard, measures of *descriptional complexity* [5] of functions have been proposed over the years, which quantify the level of complexity of an individual function, based on its shortest description. Although fundamental measures of descriptional complexity such as Kolmogorov Complexity [9] or Solomonoff Probability [10] are uncomputable, computable approximations to them have been developed [11]. In spite of this early work, there is a lack of concrete theoretical or empirical work to continue this line of investigation after [8]. Mainly, there is a lack of work that explores the relevance of such descriptional complexity measures to the generalizational ability of functions. That is, until recently, when interest in the descriptional complexity of a function was renewed as it was found that outputs of *random* maps tend to be biased towards simpler functions [12]. This result indirectly hinted that for classification tasks, the ground truth labelling function has a simple description with high probability. This was further investigated in [5] in the context of deep neural networks, where it was found that the Lempel-Ziv (*LZ*) complexity (a form of descriptional complexity) of most neural networks with random choices of weights are low (shorter description), in accordance with the result in [12]. Empirically, it was found that network weight configurations which lead to functions of smaller Lempel-Ziv complexity show better generalization performance.

These results and observations show that descriptional complexity measures may help to understand empirically the generalization behaviour of deep neural networks. In this paper, we advance this line of investigation by proposing a new theoretical framework for relating descriptional complexity measures to generalization performance. Subsequently, we also provide a computable method which improves generalization performance of neural networks by lowering their descriptional complexity.

2 Contributions

This paper makes the following contributions.

1. First, we undertake a brief theoretical analysis for exploring the relevance of such descriptional complexity measures of functions to their expected generalization error. We propose a novel measure of complexity called *Kolmogorov Growth* (KG). Error bounds are estimated which portray the dependence of generalization error of a single function f to $KG(f)$. The bounds depict that functions of higher $KG(f)$ will likely lead to a higher generalization gap. This result formalizes the Occam's razor principle for classifiers, and also concurs with the empirical findings in [5], where functions of higher *LZ*-Complexity showed worse generalization performance.

2. Like Kolmogorov Complexity, KG is also uncomputable. Therefore, we propose computable approximations of KG for neural networks, based on the concept of teacher-student approximation (similar to knowledge distillation [13]). Specifically, we show that neural network functions f which can be approximated well by smaller networks will have smaller $KG(f)$ with high probability.

3. Next, using this idea, we then develop a novel method for regularizing neural networks, called network-to-network (N2N) regularization. N2N regularization forces trained network configurations to be of low KG. We find that doing so not only improves the generalization performance across a range of training data sizes, but also helps in the case of label noise. For instance in MNIST, we see test error decrease by 94%, reaching state-of-the-art results.

4. Our study of the proposed approximation to $KG(f)$ leads to very intriguing empirical findings. For instance, we observe that typically networks show a sharp decrease in their KG as training progresses, a trend which completely reverses when the training data labels are noise corrupted.

3 Kolmogorov Growth: Relevance to Generalization

Assume that our training data consists of m number of d-dimensional i.i.d samples and their labels $S = [(z_1, y_1), (z_2, y_2), .., (z_m, y_m)]$ drawn from some distribution P^m.

We propose two growth measures for a single function f, namely Kolmogorov Growth $KG_m(f)$ and empirical Kolmogorov Growth $\widehat{KG}_S(f)$. Note that, when appropriate, we may drop the subscript m and refer to $KG_m(f)$ as $KG(f)$. These measures are primarily motivated from the well known growth function in statistics. The growth function $\Pi_m(\mathcal{F})$ is defined as $\Pi_m(\mathcal{F}) = \max_{f,\mathbf{z}\sim P^m} |\{f(z_k) : k = 1, 2, ..., m, f \in \mathcal{F}\}|$. Note that the growth function is defined for a space of functions and captures the maximum number of label assignments a function space $\mathcal{F}$ can generate on any m points in $\mathbb{R}^d$.

To define Kolmogorov Growth measures for a single function f, we first need to generate a function space from f, based on a description D_f of f. This is done via noting all the parameters involved in the description D_f (other than the co-ordinates in $\mathbb{R}^d$), and then assigning all possible values to those parameters to generate a space of functions $\mathcal{F}(D_f)$ from the description D_f. Next, we note that for a function f, we consider multiple descriptions $(D_f^1, D_f^2, D_f^3, ...)$, all of which faithfully generate the output of f over all points in $\mathcal{R}^d$, with the additional constraint that each $\mathcal{F}(D_f^i)$ should fit any m datapoints sampled from P^m.

Given these observations and denominations, we then can define the Kolmogorov Growth of a function f as follows:

$$KG_m(f) = \min_i \frac{\log \Pi_m \left(\mathcal{F}(D_f^i)\right)}{m}. \tag{1}$$

Note that $KG_m(f)$ requires the knowledge of distribution P^m of the training samples. For an instance of the datapoints in S, we define the empirical Kolmogorov Growth via the empirical growth function $\widehat{\Pi}_S(\mathcal{F})$, which only computes the number of label assignments a function space can generate over the given m training data samples in S. Thus, $\widehat{\Pi}_S(\mathcal{F}) = \max_f |\{f(z_k) : k = 1, 2, ..., m, f \in \mathcal{F}\}|$. This leads to the following definition of the empirical Kolmogorov Growth of a function f:

$$\widehat{KG}_S(f) = \min_i \frac{\log \widehat{\Pi}_S \left(\mathcal{F}(D_f^i)\right)}{m}. \tag{2}$$

Remark: Kolmorogov growth is indirectly motivated from Kolmogorov complexity itself. However, unlike Kolmogorov complexity, which is the length of the shortest program that generates f, Kolmogorov Growth is concerned with the "smallest function space" that f can belong to, that can still fit the data well. Functions which have shorter descriptions usually require a smaller number of variables and are expected to have lower Kolmogorov Growth. Moreover, it turns out that Kolmogorov Growth allows us to directly comment on the error bounds for the function f (see Section 3.1).

Remark: Note that, in the binary classification scenario, for a completely unstructured function f (i.e., f outputs random labels at every point $X \in \mathbb{R}^d$), one expects $KG(f)$ to be near its maximum value (i.e., $\log 2$). A structured f would generate shorter descriptions with less parameters and therefore lead to smaller $KG(f)$.

3.1 Bounding Generalization Error using Kolmogorov Growth

Here we present error bounds that depend on $KG(f)$, where f is the classification function, given the training data and the labels in S. First, we define the ground truth label generating function on $\mathbb{R}^d$ as g_{true}. Note that we only know the output of g_{true} on the data points in S.

We now define a set of error functions for computing training loss (0-1 loss) on S, denoted as $\widehat{err}_S(f, g_{true})$, and the overall generalization error with respect to the distribution P and the true label generating function g_{true}, denoted as $err_P(f, g_{true})$. We define them as follows:

$$\widehat{err}_S(f, g_{true}) = \sum_{i=1}^{m} \frac{(1 - f(z_i) g_{true}(z_i))}{2m} \tag{3}$$

$$err_P(f, g_{true}) = \lim_{N \to \infty, z_i \sim P} \sum_{i=1}^{N} \frac{(1 - f(z_i) g_{true}(z_i))}{2N}. \tag{4}$$

These definitions hold for all $f \in \mathcal{F}$. Note that the error functions depend both on the function f and the function g_{true}.

With this, we have the following results. The proofs of all results are provided in the supplementary material.

Theorem 1 *For $0 < \delta < 1$, with probability $p \geq 1 - \delta$ over the draw of S, we have*

$$err_P(f) \leq \widehat{err}_S(f) + \sqrt{2KG_m(f)} + \sqrt{\frac{\log(1/\delta)}{2m}}. \tag{5}$$

The following corollary of the above theorem gives bounds that depend on empirical Kolmogorov growth $\widehat{KG}_S(f)$.

Corollary 1.1 *For $0 < \delta < 1$, with probability $p \geq 1 - \delta$ over the draw of S, we have*

$$err_P(f) \leq \widehat{err}_S(f) + \sqrt{2\widehat{KG}_S(f)} + 4\sqrt{\frac{2\log(4/\delta)}{m}}. \tag{6}$$

Remark: Theorem 1 and its corollary essentially state that for functions f of lower Kolmogorov growth, we should expect a smaller generalization gap. In what follows, we outline ways to approximate the empirical Kolmorogov growth $\widehat{KG}_S(f)$.

4 Teacher-Student Approximation Bounds for Kolmogorov Growth

The fundamental idea for approximating empirical Kolmogorov growth of a function f which belongs to the function space $\mathcal{F}$ is to use a *student* classifier with a function space $\mathcal{F}^1_{small}$ (with a much smaller parametric count and complexity) to approximate the given function f (the *teacher*). We apply this idea recursively. That is, if the function $f^1_{small} \in \mathcal{F}_{small}$ approximates f best, we recursively estimate the empirical Kolmogorov growth of f^1_{small} by approximating it via another classifier with a smaller function space $\mathcal{F}^2_{small}$ (thus, $\Pi_m(\mathcal{F}^2_{small}) < \Pi_m(\mathcal{F}^1_{small})$), and so on. We use this

recursive way to then obtain a final estimate for $\widehat{KG}_S(f)$. We conjecture that, like Kolmogorov complexity itself, the true $\widehat{KG}_S(f)$ is uncomputable, so the estimate that results from this recursive approximation process is essentially an upper bound to the true $\widehat{KG}_S(f)$. The following theorem establishes an upper bound to empirical KG approximation from a single smaller student classifier.

Theorem 2 *Given the function $f \in \mathcal{F} : \mathbb{R}^d \rightarrow \mathbb{R}^2$ which outputs class logits for binary classification. We construct a function space $\mathcal{F}^1_{small}$ such that $\Pi_m(\mathcal{F}^1_{small}) < \Pi_m(\mathcal{F})$. We approximate f via another function $f^1_{small} \in \mathcal{F}^1_{small} : \mathbb{R}^d \rightarrow \mathbb{R}^2$ and let*

$$\epsilon_{max} = \max_{X \in \mathbb{R}^d} \|f^1_{small}(X) - f(X)\|^2. \tag{7}$$

Denote the output probabilities generated from the corresponding logit outputs of $f(X)$ using the softmax operator (temperature $T = 1$), as $P_0(f(X))$ (label 1 output) and $P_1(f(X))$ (label 2 output). Let $0 \leq \delta \leq 1$ be such that

$$Pr\left(\left|log\left(\frac{P_0(f(X))}{P_1(f(X)}\right)\right| \leq \epsilon_{max}\right) \leq \delta, \tag{8}$$

Then, with probability $p \geq 1 - \delta$ over the draw of S, we have

$$\widehat{KG}_S(f) \leq \delta \log 2 + \frac{\log \Pi_m\left(\mathcal{F}^1_{small}\right)}{m}, \tag{9}$$

where m is the number of samples in S.

Remark: Theorem 2 demonstrates a way to bound the true empirical Kolmogorov growth of the function f, using a single student classifier function $f^1_{small} \in \mathcal{F}^1_{small}$. The proof of Theorem 2 and the result for the case of recursive approximation are given in the supplementary material.

5 Network-to-Network (N2N) Regularization

We denote the base network to be trained as N^{base} and the function modelled by the network weights w^{base} as $N^{base}(w^{base}, X)$, where $X \in \mathbb{R}^d$ is the input. Here, $N^{base}(w^{base}, X)$ represents the output logits for the network N^{base} when presented with the input X. Thus, we have $N^{base}(w^{base}, X) \in \mathbb{R}^c$, where c is the number of classes. For what follows, let us denote the available training data and their labels by $S = \{X_i, y_i\}_{i=1}^m$.

The approach that follows is directly motivated from the result in Theorem 2. Thus, the objective of the approach is to ensure that the *KG* of the network stays low during learning, using the teacher-student approximation error in Theorem 2. This is primarily achieved by ensuring that during training, the base network function $N^{base}(w^{base}, X)$ stays near to some function modelled by a smaller network. In what follows, we outline the details of the proposed multi-level network-to-network (N2N) regularization approach.

5.1 Multi-Level N2N: Details

In multi-level N2N regularization, we have multiple smaller networks $n_1^{small}, n_2^{small}, ..., n_K^{small}$ of decreasing complexity and parametric count ($\Pi_m(\mathcal{F}^1_{small}) > \Pi_m(\mathcal{F}^2_{small}) > \cdots > \Pi_m(\mathcal{F}^K_{small})$). We also denote the corresponding function modelled by the network weight choices $w_1, w_2, .., w_m$ as $n_1^{small}(w_1, X), n_2^{small}(w_2, X), ..., n_m^{small}(w_K, X)$. Next, we outline the loss functions for all networks. For the larger to-be-trained base network N^{base}, the loss objective is to minimize cross-entropy loss on S while being close to $n_1^{small}(w_1, X)$ for some choice of weights w_1 ($\mathcal{L}_1$ in algorithm 1). For the smaller network n_1^{small}, the objective is two-fold: find the weight configuration w_1 that approximates the larger network function N^{base}, while also being close to $n_2^{small}(w_2, X)$ for some choice of w_2 ($\mathcal{L}_2$ in algorithm 1). Thus, we force the smaller network n_1^{small} to be close to the base network and an even smaller network n_2^{small} at the same time. We continue to define the loss functions for the remaining smaller networks in the same manner ($\mathcal{L}_3, \mathcal{L}_3, .., \mathcal{L}_{K-1}$), except for the final loss function $\mathcal{L}_K$ for the smallest network. The loss objective for smallest network n_K^{small} only contains a single term which enforces a weight configuration w_K which yields a network function closest to n_{K-1}^{small}. Finally, we optimize the loss functions in an alternating fashion, in the order of $\mathcal{L}_1, \mathcal{L}_2, .., \mathcal{L}_K$. Details are given in Algorithm 1. The choice of loss functions (mean-squared error based) in multi-level N2N directly follows from the result obtained in Theorem 2.

Algorithm 1 N2N Regularization (Multi-Level)

Input: Training data $\{X_i, y_i\}_{i=1}^{k}$, base network N^{base} with network weights w^{base}, K smaller networks $n_1^{small}, n_2^{small}, ..., n_K^{small}$, with network weights $w_1, w_2, ..., w_K$ (s.t. $|n_1^{small}| > |n_2^{small}| > ..|n_K^{small}|$ in size), Number of epochs J, Hyperparameters $\lambda_0, \lambda_1, \lambda_2, ..., \lambda_{K-1}$, e_{base}, e_{small}.

for $j = 1, 2, \ldots, J$ **do**
  for $i = 1, 2, \ldots, e_{base}$ **do**
    Update w^{base} using $w^{base} = w^{base} - \frac{\partial \mathcal{L}_1}{\partial w^{base}}$
    Where $\mathcal{L}_1 = \sum_{i=1}^{k}(L_{CE}(N^{base}(X_i), y_i) + \lambda_0 \|N^{base}(X_i) - n_1^{small}(X_i)\|^2)$
  for $p = 1, 2, \ldots, K$ **do**
    for $i = 1, 2, \ldots, e_{small}$ **do**
      Update w_p using $w_p = w_p - \frac{1}{k}\frac{\partial \mathcal{L}_p}{\partial w_1}$
      if $p = 1$ **then**
        $\mathcal{L}_p = \sum_i \|N^{base}(X_i) - n_1^{small}(X_i)\|^2 + \lambda_1 \|N^{base}(X_i) - n_1^{small}(X_i)\|^2$
      else if $p = K$ **then**
        $\mathcal{L}_p = \sum_i \|n_p^{small}(X_i) - n_{p-1}^{small}(X_i)\|^2$
      else
        $\mathcal{L}_p = \sum_i \|n_p^{small}(X_i) - n_{p-1}^{small}(X_i)\|^2 + \lambda_p \|n_p^{small}(X_i) - n_{p-1}^{small}(X_i)\|^2$

5.2 Other Relevant Approaches in Literature

To the best of our knowledge, our proposed approach is novel, and we did not find much relevant work. Conceptually, we found the reverse knowledge distillation method [14] to be the most relevant to our proposed approach. Reverse knowledge distillation regularizes large teacher networks using smaller, trained versions of student networks of less depth. After the student network is trained, its logit outputs are then essentially re-used for "smoothing" the outputs of the larger neural network. Here, we do not directly use trained student networks to supervise the teacher, but instead simply ensure that the teacher network while training is always within reach of *some* student network (which may change throughout the training process), which is a more relaxed constraint. We also use a mean-squared error based loss to quantify the difference between the student and the teacher network function which is motivated from Theorem 2, as opposed to the KL divergence metric used in knowledge distillation. Lastly, unlike with distillation based methods, we also use a multi-level approach for a recursive way of regularizing multiple networks of different levels of complexity.

6 Experiments

We now present the results from using our proposed N2N regularization approach on three datasets: MNIST [15], CIFAR-10 [16] and CIFAR-100 [17]. Later, we also show how N2N regularization can help with label noise on MNIST and CIFAR-10. Lastly, we present analysis of Kolmogorov growth of networks during training, with and without N2N regularization. Experiments were either carried out on an RTX 2060 GPU or a Tesla V100 GPU.

6.1 Standard Classification: MNIST, CIFAR-10, CIFAR-100

The primary objective of the experiments presented here is to see whether N2N regularization can drive the training process towards network configurations that generalize better. For each dataset, results are reported for various choices of training data size. Furthermore, to show that our regularization approach complements other commonly used regularization approaches, we show results when our approach is combined with Dropout and L2-Norm regularization. All networks were trained for a total of 200 epochs, and in each case results reported are averaged over five networks. For all experiments we set $e_{base} = 3$, $e_{small} = 1$ in Algorithm 1. The values of the regularization parameters (λ_0, λ_1) are provided in the supplementary material. Across all three datasets, we generally find that larger training data sizes need smaller regularization parameters.

Results are shown in Table 1. We note that the use of N2N regularization improves test accuracy. Mainly, we see that N2N regularization complements common regularization approaches such as dropout and L2-norm well. In all cases we find that combining these well-known regularization

MNIST					CIFAR-10				CIFAR-100			
Data Size (Training)	none	drop+l2	drop+l2 +N2N-1	drop + l2 + N2N-2	none	drop+l2	drop+l2 +N2N-1	drop + l2 + N2N-2	none	drop+l2	drop+l2 +N2N-1	drop + l2 + N2N-2
1000	96.74±0.11	97.2±0.05	97.42±0.12	**97.62**±0.08	45.35±1.25	48.99±0.71	50.58±0.65	**51.19**±0.44	6.21±0.18	7.97±0.12	8.96±0.23	**9.45**±0.2
2000	97.71±0.09	98.06±0.03	98.16±0.07	**98.27**±0.03	51.88±0.47	55.42±0.4	56.24±0.85	**57.23**±0.7	8.84±0.15	9.75±0.1	13.01±0.2	**13.75**±0.16
10000	98.92±0.04	99±0.01	99.07±0.01	**99.13**±0.01	71.84±0.29	72.13±0.29	72.55±0.31	**72.89**±0.38	33.51±0.09	35.03±0.05	35.57±0.12	**35.73**±0.04
Complete	99.38±0.04	99.59±0.03	99.68±0.04	**99.70**±0.02	92.53±0.07	92.96±0.04	**93.35**±0.05	93.26±0.09	75.65±0.1	76.32±0.07	76.65±0.1	**76.83**±0.02

Table 1: Test Accuracy on MNIST, CIFAR-10, and CIFAR-100 for different training data sizes and different regularization choices (dropout, L2-norm, single-level N2N, multi-level N2N and their combinations). We note that N2N regularization combined with dropout and L2-norm leads to improvements in test accuracy for all cases. This shows that N2N regularization complements other regularization approaches well, improving generalization performance for trained networks further. The network configurations are given in the supplementary material.

MNIST						
Noise level	F-Correction [18]	Decoupling [19]	MentorNet [20]	Co-Teaching [21]	SCE[1] [22]	N2N (2-Levels)
$p = 0.2$	98.80$\pm$0.12 (94.05)	95.70$\pm$0.02 (94.05)	96.70$\pm$0.22 (94.05)	97.25$\pm$0.03 (94.05)	98.61$\pm$0.04 (93.57)	**98.91**$\pm$0.15 (93.57)
$p = 0.5$	79.61$\pm$1.96 (66.05)	81.15$\pm$0.03 (66.05)	90.05$\pm$0.30 (66.05)	91.32$\pm$0.06 (66.05)	97.65$\pm$0.05 (65.63)	**98.05**$\pm$0.26 (65.63)
CIFAR-10						
$p = 0.2$	84.55$\pm$0.16 (76.25)	80.44$\pm$0.05 (76.25)	80.76$\pm$0.36 (76.25)	82.32$\pm$0.07 (76.25)	86.76$\pm$0.05 (80.4)	**87.55**$\pm$0.17 (80.4)
$p = 0.5$	59.83$\pm$0.17 (48.87)	51.49$\pm$0.08 (48.87)	71.10$\pm$0.48 (48.87)	74.02$\pm$0.48 (48.87)	**82.31**$\pm$0.04(46.35)	76.75$\pm$0.49 (46.35)

Table 2: Table showing the test performance on MNIST and CIFAR-10 when symmetric label noise of probabilities $p = 0.5$ and $p = 0.2$ was applied on the training data labels. Accuracies are reported for various state-of-the-art and benchmark approaches that have been recently proposed in literature, including the current state-of-the-art in symmetric cross entropy (SCE,[22]). The numbers within the brackets represent the test accuracy of the corresponding network architecture when trained with standard cross-entropy loss. We note that N2N regularization yields best performance in three of the four cases.

approaches with the proposed approach yields the best results. Furthermore, we also see that the improvements in performance do not diminish when the training data size is increased. Lastly, in most cases, we see that 2-level N2N regularization (N2N-2, $m = 2$ in Algorithm 1), outperforms single-level N2N (N2N-1, $m = 1$ in Algorithm 1), with the exception of CIFAR-10 with the full training dataset. For the CIFAR-10 and CIFAR-100 datasets, we used the benchmark ResNet architectures ResNet-44 and ResNet-50 respectively. Our results with L2-norm + dropout regularization for the ResNet-44 and ResNet-50 architectures are slightly better than the results originally reported in [23]. For the MNIST dataset, we used a 5-layer CNN with 3 conv layers and 2 fc layers. Network architecture details are provided in the supplementary material. Note that although better results can be found in literature, our objective was to demonstrate that using N2N regularization in conjunction with common regularization approaches can benefit both shallow CNN architectures (MNIST) and ResNets (CIFAR-10, CIFAR-100).

6.2 Learning with Noisy Labels: MNIST, CIFAR-10

As our proposed regularization approach constrains the network function to be simpler by minimizing an approximation of Kolmogorov Growth, it naturally applies to the case of noisy training labels. Without any regularization, label noise in the training data usually forces a network to emulate a more *complex* function, as it potentially makes the decision boundary more complex, a fact that we also empirically observe in section 6.3. We stipulate that N2N regularization should help the network in achieving simpler functions to approximate the training data labels, favoring "simpler" decision boundaries over "complex" ones, and thus potentially shielding against the corrupted labels to a certain extent. As we can simply control the level of regularization via setting the parameters $\lambda_0, ...\lambda_m$, we test whether enforcing a simpler function (large $\lambda_0, ...\lambda_m$) at the cost of compromising training loss can help improve test accuracy, when the training data is corrupted by label noise. We only tested the case where symmetric label noise of some probability p was applied (same as in [22]), and show our results for symmetric noise with $p = 0.5$ and $p = 0.2$.

First, we show the results for symmetric label noise of probability $p = 0.5$ and $p = 0.2$ on MNIST and CIFAR-10 in Table 2. Like in [21], we report the average accuracy over the last ten epochs of training. The values of the regularization parameters λ_0, λ_1 are provided in the supplementary

[1] Our implementation, using the same network architecture for a fair comparison, and optimized hyperparameter values.

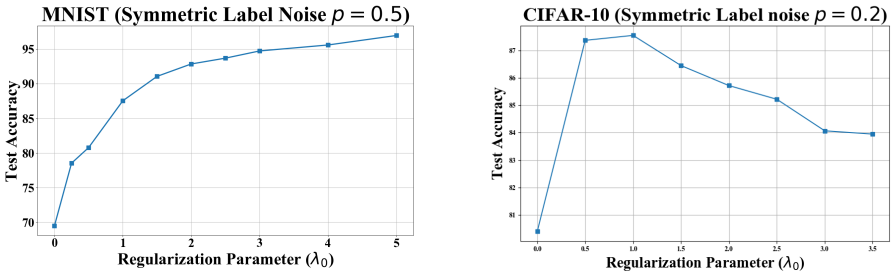

(a) (b)

Figure 1: The effect of the hyperparameter λ_0 of our proposed N2N regularization approach on test accuracy on the MNIST (symmetric noise $p = 0.5$) and CIFAR-10 (symmetric noise $p = 0.2$) datasets. The networks were regularized using 2-level N2N, with a fixed $\lambda_1 = 0.1$. In both datasets, we see significant performance improvements in response to regularization.

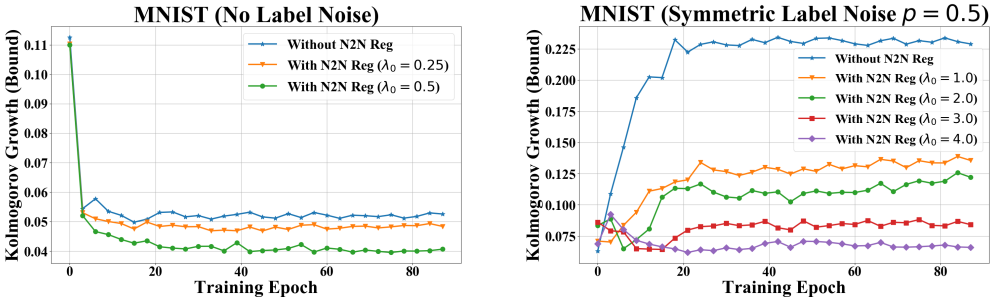

(a) (b)

Figure 2: We show the trend of the KG of networks (approximated using Theorem 2) during the training process on MNIST-1k (with and without label noise). In (a), we plot how the KG of networks changes during training on MNIST (no label noise). There, we compare the trend for three different cases: no N2N regularization, N2N regularization with $\lambda_0 = 0.25$, and N2N regularization with $\lambda_0 = 0.50$. As λ_0 increases, we see that the KG of the trained networks show a clear decrease, with the trajectories deviating after a few epochs of training. In (b), we show KG of networks during training on the MNIST dataset corrupted with symmetric label noise of probability $p = 0.5$. There, we find that the networks trained without any N2N regularization have significantly higher KG than the networks trained with N2N regularization. Particularly, we find that as λ_0 increases, KG reduces, while also improving test performance (see Fig. 1). All values of KG estimated here are 10-fold averages over multiple training sessions with randomly initialized networks.

material. We find that N2N regularization results in the best test performance in 3 out of the four cases tested, and in all cases reports quite competitive results to other benchmark methods for label noise. Next, we plot the test accuracy as a function of the regularization parameter λ_0 ($\lambda_1 = 0.1$) in Figure 1. We find that for MNIST, large λ_0 helps achieve significantly higher test accuracy, whereas for CIFAR-10, accuracy peaks around $\lambda_0 = 1$.

6.3 Comparing Kolmogorov Growth Trajectories during Training

The improvements observed via the use of N2N regularization lead to the question of how the network trajectories differ when N2N regularization is used, as compared to when it is not applied. We use the result in Theorem 2 to compute the bounded approximation to empirical Kolmogorov growth of the network function. We thus plot the approximation of $\widehat{KG}(f)$ of the function f represented by the neural network during the training process. Results are shown in Figure 2. The following summarizes our observations from the figure.

1. In the case of no training label noise, the KG of networks typically have high initial values, steeply reducing within a few epochs of training, after which it stabilizes. Expectedly, when trained with

N2N regularization, we find that the final KG of networks are lower, compared to KG of networks trained without N2N regularization.

2. In the case of label noise, we report some interesting observations. First, we see that, differently from before, the KG values rather increase with training and stabilize eventually at higher values. This is almost an opposite trend to the case of no label noise. This can be partly explained by the fact that as training progresses, the network slowly adapts its decision boundary to fit the erroneous labelling, eventually resulting in a decision boundary of high complexity.

3. For the label noise case, we find that N2N regularization significantly reduces the increase of KG during the training process. Furthermore, larger values of the λ parameter leads to networks which exhibit smaller KG values. This also helps explain the significant gains in test accuracy observed for MNIST earlier in Table 2, when using N2N regularization.

7 Discussions and Reflections

This results in this paper further the recent work by [5], where it was shown that neural networks are inherently biased towards *simpler* functions of lower Kolmogorov complexity. In particular, this paper provides an actionable solution for incorporating a function complexity prior while learning, using a novel measure called Kolmogorov Growth. Unlike Kolmogorov complexity, which is the description of the shortest program that generates some function f, Kolmogorov Growth is concerned with the "smallest function space" that f can belong to, that can still fit the data well. Functions with shorter descriptions will typically need fewer variables and thus have lower Kolmogorov Growth.

We show via our proposed N2N regularization approach, that enforcing low KG for large networks can further improve their ability to generalize. The proposed approach greatly helps in the scenario where the training data has noisy labels, attaining state-of-the-art or competitive performance on MNIST and CIFAR-10 when the training labels are corrupted with symmetric label noise. In the case of label noise, we see that networks trained without N2N regularization have large Kolmogorov Growth, which reduces immediately following the application of N2N regularization.

Our theoretical results in Section 3.1 show that networks configurations that can be approximated well by smaller networks of lower complexity will have low Kolmogorov growth, and subsequently, lower generalization error. These results concur with the very recent theoretical findings in [24], which finds analogous results for the Rademacher complexity based generalization error framework, in the context of knowledge distillation. Our main result in Theorem 1 outlines an Occam's razor like principle for generalization. Theorem 1 implies that for all functions which have zero training error, the function with the smallest $KG_m(f)$ will be the most likely to show the least generalization error.

Our empirical findings consistently show that driving the networks towards simpler functions of lower Kolmogorov Growth leads to networks that generalize better. Primarily, our proposed regularization approach ensures that during training, the network always behaves similarly to a network of lower complexity, and with multi-level N2N we apply this idea recursively. In addition to being a method for regularizing neural networks, we see that enforcing low $KG_m(f)$ on the classification function f, by increasing the regularization parameters, greatly improves performance when the training data is corrupted by label noise (Section 6.2). This also implies that a current limitation of our approach is that the hyperparameters $(\lambda_0, \lambda_1, ..)$ have to be manually tuned. Estimating their optimal values is an avenue for future research.

We observe that in N2N regularization, the properties of the smaller networks can dictate how the larger network learns. For instance, if we choose smaller networks which are highly rotation invariant in their structure (for e.g., by using a rotation-invariant CNN), then we should expect the larger networks to adopt some of the rotation invariance properties as well. To demonstrate this, we conducted an additional experiment on a custom MNIST [15] dataset, which contains images of digits translated randomly within the image. We added symmetric noise on the labels ($p = 0.5$), and tested our proposed N2N regularization approach with a student network which is highly translation invariant (large max-pooling windows). We found that in this case, N2N regularization significantly outperforms state-of-the-art methods such as SCE [22], reducing test error by 27%. This shows that it is possible to extend this work by analyzing the effect of invariance/equivariance choices in the smaller networks on the generalization behaviour of the larger network.

Finally, since our work provides a certain level of robustness against label noise, it supports activities such as crowdsourcing data labelling, which potentially contains significant label noise.

References

[1] D.H. Wolpert and W.G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, 1997.

[2] Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, March 2003.

[3] Anselm Blumer, A. Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *J. ACM*, 36(4):929–965, October 1989.

[4] Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Localized rademacher complexities. In Jyrki Kivinen and Robert H. Sloan, editors, *Computational Learning Theory*, pages 44–58. Springer, 2002.

[5] Guillermo Valle-Perez, Chico Q. Camargo, and Ard A. Louis. Deep learning generalizes because the parameter-function map is biased towards simple functions. In *International Conference on Learning Representations*, 2019.

[6] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. 2017. Published in ICLR 2017.

[7] Occam's razor. `https://en.wikipedia.org/wiki/Occam%27s_razor`.

[8] Jürgen Schmidhuber. Discovering neural nets with low kolmogorov complexity and high generalization capability. *Neural Networks*, 10(5):857–873, 1997.

[9] A.N. Kolmogorov. On tables of random numbers. *Theoretical Computer Science*, 207(2):387–395, 1998.

[10] R.J. Solomonoff. A formal theory of inductive inference. part i. *Information and Control*, 7(1):1–22, 1964.

[11] A. Lempel and J. Ziv. On the complexity of finite sequences. *IEEE Transactions on Information Theory*, 22(1):75–81, 1976.

[12] Kamaludin Dingle, Chico Q. Camargo, and Ard A. Louis. Input–output maps are strongly biased towards simple outputs. *Nature Communications*, 9(1):761, Feb 2018.

[13] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network. *arXiv e-prints*, March 2015.

[14] Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[15] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. `http://yann.lecun.com/exdb/mnist/`, 2010.

[16] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research).

[17] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-100 (canadian institute for advanced research).

[18] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2233–2241, 2017.

[19] Eran Malach and Shai Shalev-Shwartz. Decoupling "when to update" from "how to update". In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[20] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *ICML*, 2018.

[21] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, pages 8535–8545, 2018.

[22] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *IEEE International Conference on Computer Vision*, 2019.

[23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society, 2016.

[24] Daniel Hsu, Ziwei Ji, Matus Telgarsky, and Lan Wang. Generalization bounds via distillation. In *International Conference on Learning Representations*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] **We clearly outline the motivation and contributions of our approach in the abstract.**
   (b) Did you describe the limitations of your work? [Yes] **In section 7, we describe a limitation of our approach.**
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] **Our work is largely theoretical and has minimum societal impact. One aspect that the work could impact is crowdsourcing data labelling, which potentially contains significant label noise. Since the work provides a certain level of robustness against label noise, it could encourage more crowdsourcing.**
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] **We are providing the proofs of all results in the supplementary material.**
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] **Code and instructions will be provided in the supplementary material.**
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] **We report the standard deviation of accuracies for our main experiments (Section 6.1 and 6.2)**
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A] **Datasets are open-source**
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] **Code will be included in the supplemental material**
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]