

# United We Stand: Using Epoch-Wise Agreement of Ensembles to Combat Overfit\*

Uri Stern, Daniel Schwartz, Daphna Weinshall

School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel  
ustern@gmail.com, Daniel.Schwartz1@mail.huji.ac.il, daphna@mail.huji.ac.il

## Abstract

Deep neural networks have become the method of choice for solving many classification tasks, largely because they can fit very complex functions defined over raw data. The downside of such powerful learners is the danger of overfit. In this paper, we introduce a novel ensemble classifier for deep networks that effectively overcomes overfitting by combining models generated at specific intermediate epochs during training. Our method allows for the incorporation of useful knowledge obtained by the models during the overfitting phase without deterioration of the general performance, which is usually missed when early stopping is used.

To motivate this approach, we begin with the theoretical analysis of a regression model, whose prediction – that the variance among classifiers increases when overfit occurs – is demonstrated empirically in deep networks in common use. Guided by these results, we construct a new ensemble-based prediction method, where the prediction is determined by the class that attains the most consensual prediction throughout the training epochs. Using multiple image and text classification datasets, we show that when regular ensembles suffer from overfit, our method eliminates the harmful reduction in generalization due to overfit, and often even surpasses the performance obtained by early stopping. Our method is easy to implement and can be integrated with any training scheme and architecture, without additional prior knowledge beyond the training set. It is thus a practical and useful tool to overcome overfit.

## 1 Introduction

Deep supervised learning has achieved exceptional results in various image recognition tasks in recent years. This impressive success is largely attributed to two factors – ready access to very large annotated datasets, and powerful architectures involving a huge number of parameters, thus capable of learning very complex functions from large datasets. However, when training very large models, we face the risk of *overfit* – when models are fine-tuned to irrelevant details in the training set. It is defined here as the co-occurrence of *increase* in training accuracy and *decrease* in test accuracy, which implies poorer generalization as training pro-

\*This work was supported by grants from the Israeli Council of Higher Education and the Gatsby Charitable Foundation.  
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An illustrative example, showing the hypothetical behavior of an ensemble of size 3. (a) Visualization of each model’s prediction over time (5 epochs) (b) The final prediction of the 3 models (top row), regular ensemble (bottom left), and our method (bottom right).

ceeds. Not surprisingly, given the gravity of this issue, a variety of measures have been developed through the years to combat overfit, as reviewed in Section 2.

One important cause of overfit in deep learning is the existence of false labels in the dataset (also termed “noisy labels”), which the network nevertheless is able to memorize. Empirically, such noisy labels are known to be memorized late by deep networks (Arpit et al. 2017), resulting in overfit - errors that occur later in training. As deep learning requires large annotated datasets, the problem of false labels often arises in real applications, medical data being an important example (Karimi et al. 2020), where the reliable labeling of data is expensive. Moreover, cheap alternatives such as crowd-sourcing or automatic labeling typically introduce noisy labels to the dataset (Nicholson, Sheng, and Zhang 2016). It is thus important to develop additional tools for combating overfit, especially heavy overfit caused by noisy labels, which is our goal in this paper.

Many methods have been developed to reduce overfit (See Section 2), often by adding some form of regularization to the network’s training. One important example is *early stopping*, where the training is concluded ahead of time to avoid overfit. The problem with early stopping is that models can learn useful features from the data even as they are overfitting, which raises the need for additional methods that reduce overfit without limiting the training or the model.

Our proposed method belongs to the family of deep ensemble classifiers. Differently from most methods, our new ensemble classifier does not only consider the final predictions of the networks in the ensemble, but also tracks the networks’ responses as learning proceeds. In this context, we recall recent empirical findings (Hacohen, Choshen, and

Weinshall 2020; Pliushch et al. 2021), which show that when the predictions of networks over the train and test data are tracked through time, almost all the networks will either predict correctly the correct label or almost all will fail. These findings imply that deep networks demonstrate high agreement per epoch in their predictions of correct labels. But while all the networks are shown to succeed together, they may still fail in different ways, as visualized in Fig. 1. Here we pursue the hypothesis that **such diversity can be exploited to distinguish between false predictions and correct ones**, most effectively when overfit occurs.

Our work begins by investigating the hypothesis mentioned above, where we study networks’ agreement over false prediction both theoretically and empirically. We start with the theoretical analysis of a regression model (Section 4.1), which reinforces the intuition that *overfit increases the prediction variance between linear models*. Accordingly, we make two conjectures: (i) ensembles become more effective when overfit occurs; (ii) the agreement between networks is reduced when overfit occurs. These conjectures are verified empirically in Section 4.2 while evaluating deep networks in common use. We see that ensemble classifiers indeed improve performance when overfit occurs. However, overfit is **not** eliminated. Conjecture (ii) is validated with a *new empirical result*: the variance of correct test predictions is much smaller than the variance of false predictions.

The aforementioned empirical result implies that correct predictions at inference time can be distinguished from false ones by looking at the *variance in predictions between multiple networks throughout the training epochs*. This is used to construct a new algorithm for ensemble-based prediction, which is much more resistant to overfit (see Section 5.1). This method is then tested on different image classification and text classification datasets with label noise where the networks manifest significant overfit (see Section 5.2). In these experiments we see two positive effects of our method: (i) overfit is eliminated in almost all the experiments (except in extreme cases of label noise), showing that the method is *robust against overfit*; (ii) its performance is superior to the ensemble’s best epoch, when identified for early stopping.

To conclude, our **main contribution** is a new method for ensemble-based prediction, which is resistant to overfit, and does not require any prior knowledge on the data or changes to the training protocol. It is applicable with any network and dataset, easy to implement, and is effective *as a practical tool against overfit*. Most importantly, when the models are still able to learn useful patterns from the data after the occurrence of overfit, for example, when the test error shows “double descent” (see Stephenson and Lee 2021; Nakkiran et al. 2021; Stern and Weinshall 2023), our method allows the user to use this forgotten knowledge while reducing overfit, outperforming the optimal early stopping. Finally, our method can be readily incorporated with existing methods, which are designed to combat overfit, in order to boost their performance, as shown in Section 6. Our code is available at <https://github.com/uristern123/United-We-Stand-Using-Epoch-wise-Agreement-of-Ensembles-to-Combat-Overfit>.

## 2 Prior Work

Previous work on combating overfit can be largely divided into two groups, which are **not** mutually exclusive. **The first group**, which includes the majority of works, focuses on modifying *the learning process* in order to prevent or reduce overfit. These include various forms of data augmentation (Shorten and Khoshgoftaar 2019) and regularization techniques, such as dropout (Srivastava et al. 2014), weight decay (Krogh and Hertz 1991), batch normalization (Ioffe and Szegedy 2015) and early stopping - the termination of training before overfit occurs. Such methods come at a cost, as they (and specifically early stopping) often limit the ability of the model to learn useful features from the training data in the late stages of learning, and even in the overfit phase, wherein useful features can still be learned by the model (Stern and Weinshall 2023). Many methods have also been suggested for training under the presence of label noise, a known cause for overfit, see (Song et al. 2022) for a recent review.

**The second group**, to which our method belongs, comes into play after learning is concluded. Such methods involve the design of post-processing algorithms to be invoked during *inference time* (e.g., Lee et al. 2019). Relevant work, when dealing with label noise, includes (Zhang, Lee, and Agarwal 2021; Bae et al. 2022). An ensemble classifier is used in (Salman and Liu 2019) to combat overfit due to insufficient training data, using network confidence to filter out erroneous predictions. Our method differs primarily in that it provides labels for all points, implicitly correcting false labels of regular ensembles, which is achieved by relying on a novel analysis of the ensemble’s predictions.

One benefit of the second approach is that it doesn’t require additional time for training, e.g., by way of re-training or data-specific hyper-parameter tuning, which is a major drawback of the first approach. Additionally, methods from the second group can be incorporated into any training scheme, which is useful when methods from the first group fail to completely eliminate overfit. Lastly, methods from this group do not prevent the model from learning useful features at the late stages of training, as some of the first group methods do.

**Ensemble methods.** Engaging a set of classifiers, rather than a single one, is an established methodology to deal with errors (Hansen and Salamon 1990). Ensembles of deep neural networks are typically aggregated by simple yet useful methods. The two most widely used aggregation methods are the *majority vote* and *class probabilities average* (Lakshminarayanan, Pritzel, and Blundell 2017) of the ensemble.

Generating diversity in the ensemble can be done in a variety of ways (Li, Wang, and Ding 2018; Ganaie, Hu et al. 2021), such as training with different datasets, as in bagging (Breiman 1996) and boosting (Freund and Schapire 1997), or by varying weights and hyper-parameters (Wenzel et al. 2020). A natural cause of diversity is learning from noisy labels, as illustrated in (Shwartz, Stern, and Weinshall 2022), where ensembles of neural networks are used to identify such labels. As ensembles are expensive, some papers focus on minimizing their cost rather than maximizing their performance, see for example (Wen, Tran, and Ba 2020).

### 3 Methodology

**Notations** Let  $N$  denote the number of networks in an ensemble of deep neural networks trained independently on a training dataset with  $C$  classes for  $E$  epochs using Stochastic Gradient Descent (SGD), and tested after each epoch of training on the test set  $\mathbb{T}$ . Let  $f_i^e(x)$  denote the prediction of network  $i$  for  $\mathbf{x} \in \mathbb{T}$  in epoch  $e$ .  $N^e(x)$  denotes the ensemble prediction in each epoch  $e$ , defined as:

$$N^e(\mathbf{x}) = \operatorname{argmax}_{c \in C} \sum_{i=1}^N \mathbb{1}_{[f_i^e(\mathbf{x})=c]} \quad (1)$$

**Label noise generates overfit** As deep models can memorize any data distribution, having false labels in the training set (also termed noisy labels) leads networks to overfit, since their test accuracy decreases by the end of training when false labels are being memorized. We thus expose our models to data with label noise in order to evaluate robustness to heavy overfit. We evaluate our method on: (i) datasets with synthetic label noise, (ii) real-world datasets with label noise crafted via unreliable web-labeling (Clothing1M, Webvision), and (iii) a dataset with inherent label noise due to inherently confusing data for annotators (Animal10N).

**Injecting label noise** To generate data with synthetic label noise, we use two standard noise models (Patrini et al. 2017):

1. **Symmetric noise (SN):** a fraction  $p \in \{0.2, 0.4, 0.6\}$  of the labels is selected randomly. Each selected label is switched to any other label with equal probability.
2. **Asymmetric noise (AsN):** a fraction  $p$  of the labels is selected randomly. Each selected label is switched to another label using a deterministic permutation function.

**Inter-Model Agreement** In order to measure agreement between different models during inference time, we define below in (2) an Agreement score. This score measures the average fraction of networks in the ensemble that predict class  $c$  at point  $\mathbf{x}$  in a set of training epochs  $\mathcal{E}$ .

$$\operatorname{Agr}(\mathbf{x}, c) = \frac{1}{N|\mathcal{E}|} \sum_{e \in \mathcal{E}} \sum_{i=1}^N \mathbb{1}_{[f_i^e(\mathbf{x})=c]} \quad (2)$$

### 4 Overfit and Ensemble Classifiers

In this section we study, theoretically and empirically, the dynamics of ensembles when overfit actually occurs. We begin with a theoretical analysis of a regression model of linear classifiers (Section 4.1), showing that the agreement between such classifiers decreases when overfit occurs. In Section 4.2 we empirically study the relevance of this result to deep learning, showing that the same phenomenon occurs also with deep neural networks. We conclude in Section 4.3 with the analysis of the onset time of correct and false predictions. This empirical analysis shows that false predictions caused by overfit are associated with lower agreement scores as compared to the correct predictions, implying that false predictions are less common throughout the training than the corresponding correct predictions. Later in Section 5.1, these results guide the construction of a new ensemble-based prediction algorithm, which is resistant to overfit.

#### 4.1 Overfit in Linear Models: Theoretical Result

Since deep learning models are difficult to analyze theoretically, common practice invokes simple models (such as linear regression) that can be theoretically analyzed, whose analysis may shed light on the behavior of actual deep models. Accordingly, we formally analyze an ensemble of linear regression models trained using gradient descent via the perspective of their agreement. Our analysis culminates in a theorem, which states that the agreement between linear regression models decreases when overfit occurs in all the models and their individual generalization error increases.

**Theorem.** *Assume that all models are trained using the same training set, and all suffer from overfit at the same time. Then, under reasonable assumptions, the disagreement between the models increases at the time overfit occurs.*

Here is a brief sketch of the proof<sup>1</sup>:

1. We measure *Disagreement* by the empirical variance over models of the error vector at each test point, averaged over the test examples.
2. We prove the following (intuitive) Lemma: *Overfit occurs in a model iff the gradient step of the model, which is computed from the training set, is negatively correlated with a vector unknown to the learner - the gradient step defined by the test set.*
3. We show that under certain asymptotic assumptions, the disagreement is approximately  $-\rho$ , where  $\rho$  denotes the average correlation between each network's gradient step when computed using either the training set or the test set. It now readily follows from the Lemma that if overfit occurs in all the models then the average correlation  $\rho$  must be negative, and the aforementioned disagreement score increases.


#### 4.2 Overfit in Deep Networks: Empirical Study

We discuss in the introduction known findings concerning ensembles of deep models, which imply that all networks are likely to succeed or fail together in their prediction. However, it is still possible that they fail in different ways, where each network predicts a different false label. With overfit in particular, they may favor different false labels at different epochs (see illustration in Fig. 1).

We therefore start our empirical investigation by analyzing the variance of false predictions (or errors) in such ensembles. We use datasets injected with label noise in this study, in order to amplify the prevalence of overfit. The diversity in the ensembles has 2 sources: (i) different random initialization, and (ii) different random mini-batches within each epoch, where all networks train on the full datasets.

More specifically, we begin by training an ensemble of  $N$  DNNs on noisy labeled datasets and then isolate the test points with erroneous prediction in the last epoch. Using this set of points, we compute the *Error Consensus Score (ECS)*, which measures the number of networks that agree on each erroneous prediction (from 1 to  $N$ ). In Fig. 2 we plot some empirical histograms of *ECS*.


<sup>1</sup>The full proof is given by Stern, Shwartz, and Weinshall (2023), see Appendix B.



(a) C10 20% AsN (b) C100 60% SN (c) C100 10% AsN

Figure 2: Histograms of  $ECS$  - the number of networks that agree on each erroneous prediction (out of 5).

The results in Fig. 2 indicate that when erroneous predictions are concerned, their distribution seems to have a large variance. In fact, in many of our study cases, the minority value ( $ECS < \frac{N}{2}$ ) dominates the distribution. We also see that in more difficult settings (e.g., Cifar100), more mistakes are made by the majority, making the ensemble unable to correct them. We therefore conjecture that ensemble classifiers will be effective in correcting erroneous predictions in cases of overfit. In addition, we conjecture that with minor overfit, ensembles will only make a little difference.



(a) C10, 60% SN


(b) C100, 40% SN

Figure 3: Test accuracy over the epochs of a single network (green) and ensemble (red) trained on a dataset with label noise, as well as the test accuracy of a single network (blue) and ensemble (orange) trained on the clean data only.

These conjectures are supported by our empirical results, in experiments involving various datasets and noise settings, as shown in Fig. 3 and Table 1. In particular, Fig. 3 clearly shows that the benefit of the ensemble classifier is much larger when there is significant overfit, which is seen when the training data has significant label noise. However, Fig. 3 also shows that **while ensembles improve accuracy when there is overfit, they do not eliminate the phenomenon** - there is still performance deterioration as training proceeds. The "rebound" in test accuracy of the single network is due to the "epoch-wise double descent" phenomenon, see details in (Stephenson and Lee 2021; Nakkiran et al. 2021).

### 4.3 Remember Your History and Avoid Overfit

Ensembles are most effective when the predictions of the different members of the ensemble show a large variance, as discussed above. We now inspect another potential source of variability - a prediction's persistence. Following our theoretical and empirical analysis and the agreement score defined in (2), we hypothesize that since erroneous predictions on the test examples have large variance, correct predictions will have larger persistence (agreement) than erroneous ones, amplifying the difference between the two.



(a) C10, 60% SN (b) TImg, 40% SN (c) Img100, 40% SN

Figure 4: Histograms of the margin score (3): blue indicates correct predictions and orange erroneous predictions.

---

#### Algorithm 1: max Agreement Prediction (MAP)

---

**Input:**  $f_i^e(x)$  - the prediction of network  $i$  for test example  $\mathbf{x} \in \mathbb{T}$ ,  $\forall \mathbf{x}, i, e \in \mathcal{E}$

**Output:** final prediction  $\forall \mathbf{x} \in \mathbb{T}$

**for**  $\mathbf{x}$  in  $\mathbb{T}$  **do**

final\_prediction\_arr[ $\mathbf{x}$ ]  $\leftarrow \operatorname{argmax}_c Agr(\mathbf{x}, c)$

**end for**

Return final\_prediction\_arr;

---

In order to test this conjecture, we use the agreement score (2) defined in Section 3. More specifically, we compare the agreement score of the ensemble's final prediction in epoch  $E$  ( $N^E(x)$ ) with the agreement score of the most agreed upon label, which is different from the final prediction:

$$AgrMargin(\mathbf{x}, y) = Agr(\mathbf{x}, y) - \max_{c \neq y} Agr(\mathbf{x}, c) \quad (3)$$

$$y = N^E(\mathbf{x})$$

Fig. 4 shows histograms of score (3), separated to correct and erroneous predictions. Clearly, most of the predictions with negative margin scores are false, while the vast majority of the correct predictions have a positive score. The phenomenon is rather general, shown in a variety of conditions in which overfit is enhanced by injecting label noise.

## 5 Proposed Method (MAP)

We present a new ensemble classifier algorithm called Max Agreement Prediction (MAP) in Section 5.1, and demonstrate its superior performance in Section 5.3.

### 5.1 Algorithm

Fig. 4 shows that empirically, the agreement margin score in (3) can be reliably used to identify false predictions. We take this result one step further and propose to use the ensemble statistics of agreement score in order to select the label prediction, instead of the usual practice of selecting  $N^E(\mathbf{x})$  from (1). Specifically, we propose the following prediction selection rule:

$$y(\mathbf{x}) = \operatorname{argmax}_c Agr(\mathbf{x}, c) \quad (4)$$

We use this score in Alg. 1 and test it in Section 5.3, where its boost in performance is shown.

### 5.2 Empirical Setup

To evaluate our method with different levels of overfit we use image and text classification datasets with injected noisy

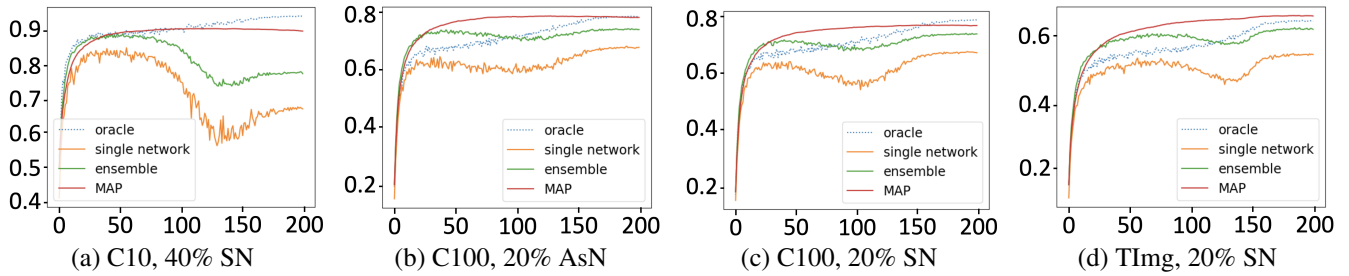


Figure 5: Mean test accuracy of our method (MAP, in red) in each training epoch (X-axis), compared with the following baselines: (i) A single network trained on the noisy dataset (orange). (ii) An ensemble trained on the noisy dataset (green). (iii) An oracle of a *single network* trained on the clean subset of the data (dotted blue). Our method shows significant *and stable* improvement, and largely eliminates the overfit caused by the noisy labels. Note that even though in some cases the test accuracy of the network/ensemble improves at the late stages of the training, it is often inferior compared to its past performance, and is *always* inferior to our method’s performance.

labels (Cifar10/100, TinyImagenet, Imagenet100, MNLI, QNLI, QQP) and datasets with native label noise (Webvision50, Clothing1M, Animal10N)<sup>2</sup>. Training involves common methods known to reduce overfit, such as data augmentation, batch normalization and weight decay, in order to capture our method’s *added value* in combatting overfit.

As comparison baseline we use the following methods: (i) A single network of the same architecture and similarly trained. (ii) The same ensemble while using the ‘majority vote’ prediction rule -  $N^E(\mathbf{x})$  defined in (1). (iii) The same ensemble while using the ‘class probabilities average’ prediction rule. The last 2 baselines are instances of a ‘regular ensemble’, reflecting methods in common use that incur the same training cost as our method. A comparison with methods of comparable *inference time* are discussed in Section 6.

Additionally, we compare our method to alternative post-processing methods, which also aim to improve classification at inference time. To assure a fair comparison, all the results are obtained using the same network architectures in the ensemble, which is most relevant to the methods that make use of an ensemble as we do (Salman and Liu 2019; Lee et al. 2019). The method described in (Bae et al. 2022) is executed several times with random seeds, the results of which are processed by the majority vote as commonly done.

Lastly, we test our method’s added value when integrated with two existing methods for training ensembles: (i) batch ensemble (Wen, Tran, and Ba 2020), which aims to reduce costs; and (ii) hyper ensemble (Wenzel et al. 2020), which aims to increase the ensemble’s diversity by using different hyper-parameters as well as different initialization.

### 5.3 Results

Fig. 5 and Table 1 summarize the test accuracy results. As performance drop becomes more severe (e.g., due to increased label noise), our method significantly outperforms the regular ensemble (both majority vote and class probabilities average) at the end of the training. It often outperforms

optimal early stopping (i.e. the ensemble’s best performance before the overfit, see Fig. 5). Lastly, Table 3 shows that our method is complementary to ensemble training methods that aim to increase diversity or decrease their cost.

Importantly, note that MAP eliminates the overfit in almost all cases: when inspecting the case studies shown in Fig. 5, we clearly see that the test accuracy does not deteriorate when MAP is used, unlike a single network and the regular ensemble, making it a superior alternative to early stopping. Only in severe, unrealistic settings of label noise (such as 40% asymmetric noise) do we see a deterioration in our method’s performance in the late stages of the training (though MAP still outperforms the regular ensemble).

### 5.4 Limitations

Our method has two main limitations: (i) the need for multiple network training, and (ii) the need to save multiple checkpoints of the network to be used at inference time. However, using multiple GPUs can mitigate these limitations, and in Section 6 and Table 3 we show that training a few networks and saving only a few checkpoints for each one, suffice for optimal or almost optimal results.

## 6 Ablation Study

Our method requires the training of a few networks and the retaining of multiple checkpoints from each one, which is computationally costly. This is justified by the large improvement achieved using our method, as shown above. In order to evaluate the practical implication of those added complexities, we investigate in Section 6.1 the effect of the number of networks in the ensemble on the final outcome of a regular ensemble and of our method, and in Section 6.2 the effect of using only a small subset of the training checkpoints. In Section 6.3 we evaluate our method on datasets in which the size of the training set is small, which is another known cause for overfit, and without label noise. In section 6.4 we check our method’s additive value when unique methods for learning with label noise are used.

Summarizing the full ablation results, even with a *few networks* and a *few checkpoints for each*, our method achieves optimal or near-optimal performance, making it practical

<sup>2</sup>Additional details are provided in (Stern, Shwartz, and Weinshall 2023): App. E for implementation details, App. C for additional evaluations, and App. D for additional ablation results.



Method/Dataset	Cifar10 sym			Cifar100 sym			Clothing1M
% label noise	20%	40%	60%	20%	40%	60%	38% (est)
<i>single network</i>	85.4 $\pm$ .1	67.7 $\pm$ .6	43.5 $\pm$ .5	67.0 $\pm$ .2	51.1 $\pm$ .2	32.2 $\pm$ .4	65.1 $\pm$ .1
<i>majority vote ensemble</i>	90.1 $\pm$ .2	78.5 $\pm$ .7	54.0 $\pm$ .4	73.5 $\pm$ .2	61.2 $\pm$ .4	42.8 $\pm$ .2	66.0 $\pm$ .2
<i>probability average ensemble</i>	90.5	79.5	56.6	74.9	64.2	46.3	67.1
<i>MAP</i>	<b>93.2 <math>\pm</math> .1</b>	<b>90.0 <math>\pm</math> .1</b>	<b>83.8 <math>\pm</math> .7</b>	<b>76.7 <math>\pm</math> .1</b>	<b>69.7 <math>\pm</math> .5</b>	<b>60.0 <math>\pm</math> .4</b>	<b>71.7 <math>\pm</math> .1</b>

Method/Dataset	TinyImagenet sym		Imagenet100 sym	
% label noise	20%	40%	20%	40%
<i>single network</i>	54.5 $\pm$ .7	40.3 $\pm$ .2	76.4 $\pm$ .5	64.9 $\pm$ .1
<i>majority vote ensemble</i>	62.0 $\pm$ .1	50.3 $\pm$ .4	82.8 $\pm$ .2	75.5 $\pm$ .5
<i>probability average ensemble</i>	64.0	53.2	83.9	77.9
<i>MAP</i>	<b>66.0 <math>\pm</math> .1</b>	<b>60.1 <math>\pm</math> .1</b>	<b>83.9 <math>\pm</math> .1</b>	<b>80.9 <math>\pm</math> .3</b>

Method/Dataset	Cifar10 asym			Cifar100 asym			Animal10N
% label noise	10%	20%	40%	10%	20%	40%	8% (est)
<i>single network</i>	90.8 $\pm$ .1	83.1 $\pm$ .2	59.7 $\pm$ .3	74.1 $\pm$ .2	67.5 $\pm$ .2	47.5 $\pm$ .1	86.0
<i>majority vote ensemble</i>	93.5 $\pm$ .1	88.4 $\pm$ .2	64.0 $\pm$ .4	79.0 $\pm$ .1	73.7 $\pm$ .1	53.4 $\pm$ .2	87.3
<i>probability average ensemble</i>	93.4	88.6	63.3	79.8	74.5	52.9	87.9
<i>MAP</i>	<b>95.2 <math>\pm</math> .1</b>	<b>94.4 <math>\pm</math> .2</b>	<b>85.6 <math>\pm</math> .2</b>	<b>80.6 <math>\pm</math> .08</b>	<b>77.7 <math>\pm</math> .1</b>	<b>57.5 <math>\pm</math> .2</b>	87.4

Method/Dataset	mnli		qnli		qqp	
% label noise	20%	40%	20%	40%	20%	40%
<i>single network</i>	79.3 $\pm$ .07	74.5 $\pm$ .06	86.0 $\pm$ .05	74.5 $\pm$ .05	85.7 $\pm$ .04	75.6 $\pm$ .04
<i>majority vote ensemble</i>	81.6 $\pm$ .2	76.7 $\pm$ .2	87.3 $\pm$ .1	74.6 $\pm$ .4	88.1 $\pm$ .1	76.3 $\pm$ .1
<i>probability average ensemble</i>	82.1	77.3	87.5	74.3	88.2	77.0
<i>MAP</i>	<b>82.6 <math>\pm</math> .09</b>	<b>79.3 <math>\pm</math> .1</b>	<b>89.0 <math>\pm</math> .1</b>	<b>82.5 <math>\pm</math> .08</b>	<b>88.8 <math>\pm</math> .05</b>	<b>82.2 <math>\pm</math> .1</b>

Table 1: Mean test accuracy (in percent) and standard error, comparing our method (MAP) and some baselines. In the top 3 tables, we show results with 4 image datasets with injected label noise and 2 image datasets with presumed label noise marked by (est). In the bottom table, we show results with 3 text datasets with injected label noise.

Method/Dataset	Cifar10 sym		Cifar100 sym		TinyImagenet sym		Clothing1M
% label noise	20%	40%	20%	40%	20%	40%	38% (est)
<i>RoG</i>	87.4	81.8	64.3	55.6	-	-	68.0
<i>consensus</i>	87.4 $\pm$ .3	74.7 $\pm$ .3	71.0 $\pm$ .4	41.9 $\pm$ .1	23.7 $\pm$ .2	18.5 $\pm$ .3	66.9
<i>NPC</i>	89.8	77.5	73.7	61.5	62.1	50.3	70.8
<i>MAP</i>	<b>93.2 <math>\pm</math> .07</b>	<b>90.0 <math>\pm</math> .1</b>	<b>76.7 <math>\pm</math> .1</b>	<b>69.7 <math>\pm</math> .5</b>	<b>66.0 <math>\pm</math> .1</b>	<b>60.1 <math>\pm</math> .06</b>	<b>71.7 <math>\pm</math> .2</b>

Table 2: Mean test accuracy (in percent) and standard error, comparing our method (MAP) to some of the post-processing methods discussed above. We report 3 image datasets with injected label noise and a single dataset with presumed annotation errors marked by (est). Source of alternative methods results: RoG from (Lee et al. 2019), consensus (Salman and Liu 2019) and NPC (Bae et al. 2022) results are recomputed using our settings and our own implementation.

Method	BatchEnsemble		Hyper-ensemble	
% label noise (sym)	20%	40%	20%	40%
<i>original accuracy</i>	69.5	52.8	72.1	61.3
<i>adding MAP</i>	<b>72.5</b>	<b>60.1</b>	<b>74.1</b>	<b>67.6</b>

Method	BatchEnsemble		Hyper-ensemble	
% label noise (asym)	10%	20%	10%	20%
<i>original accuracy</i>	75.7	67.8	76.9	71.2
<i>adding MAP</i>	<b>77.8</b>	<b>71.1</b>	<b>77.5</b>	<b>75.3</b>

Table 3: The added value of our method when integrated with 2 existing ensemble methods designed to decrease cost or increase diversity, on Cifar100 with label noise.

and useful even with limited computational resources. Our method can be combined with methods for learning with noisy labels, improving performance when such methods fail to eliminate overfit, and is also useful when the training set is very small. Finally, the emerging picture is robust to

changes in the networks’ architecture, and our method maintains its benefit over alternative deep ensemble<sup>2</sup>.




Figure 6: test accuracy of a regular ensemble and MAP when varying the number of networks in the ensemble (x-axis).

## 6.1 The Effect of Ensemble Size

In the results shown in Table 1, all ensemble classifiers use the same number of trained networks (5). However, when


considering the number of networks evaluated at inference time and in order to achieve a fair comparison between regular ensembles and our method, we repeat the study without limiting the ensemble’s size. Accordingly, we report accuracy at the point where the addition of models does not improve performance, see results in Table 4. Notably, our method maintains its superiority. Moreover, we see that with MAP - even a few networks achieve optimal or close to optimal results, making it practical for use (see Fig. 6).

Method/Dataset	Cifar10 sym		Cifar100 sym	
% label noise	20%	40%	20%	40%
MAP (5)	93.3	90.0	76.6	70.1
majority vote (opt)	91.8	84.2	77.3	69.1
MAP (opt)	93.5	90.7	77.4	71.3
Method/Dataset	Cifar10 asym		Cifar100 asym	
% label noise	20%	40%	20%	40%
MAP (5)	94.4	85.7	77.9	57.3
majority vote (opt)	90.5	66.5	76.6	56.1
MAP (opt)	94.5	86.9	78.4	59.1

Table 4: Comparing ensembles with optimal size (where more networks doesn’t improve performance), denoted (*opt*), to our method with 5 and *opt* networks.

## 6.2 How Many Checkpoints Are Needed?

The agreement score in (2) can be estimated from any subset of epochs. Fig. 7 shows test accuracy when using a subset of the training epochs, equally spaced through time, showing that even a few checkpoints are sufficient to achieve the same accuracy (or close to it) as when using all the epochs.



C10 40% AsN C100 60% SN TImg 20% SN C100 20% AsN

Figure 7: Test accuracy of MAP when using only a fraction of the epochs (equally spaced) per number of epochs.

## 6.3 Performance Evaluation on Clean Datasets




Figure 8: X-axis: number of datapoints per class. Y-axis: fraction of predictions of a regular ensemble, which MAP changes from true to false (orange) and from false to true (blue). MAP outperforms the ensemble whenever the blue curve lies above the orange curve.

We evaluate our method on clean datasets of different sizes, where overfit can be much smaller or nonexistent. Results are shown in Fig. 8. Clearly, MAP can be beneficial in such circumstances as well.

## 6.4 Combining MAP With Unique Method for Learning With Label Noise

As label noise is a major cause of overfit, we wanted to check whether our method can provide *additional* performance gain on top of unique training methods designed to improve robustness to label noise, especially when they do not fully manage to eliminate the overfit. In Table 5 we test our method on models trained with the ELR method (Liu et al. 2020), and show that our method can still provide significant added value to this method.

Method/Dataset	TinyImagenet		
% label noise (sym)	20%	40%	60%
ELR	57.5	47.8	25.9
ELR ensemble	61.6	54.0	33.3
ELR + MAP	<b>62.3</b>	<b>57.4</b>	<b>47.2</b>

Table 5: A competitive method for learning with label noise - ELR (Liu et al. 2020), and its performance with/out MAP.

## 7 Conclusions

Overfit is a notorious problem, which afflicts deep learning, especially in the context of real-life data. In this paper, we propose a new ensemble classifier based on a collection of deep networks, termed MAP, which uses the evolution of predictions in each model to generate predictions that in most cases are resistant to overfit. This result is consistently shown using various datasets, including both textual and image data, in a variety of settings that exhibit heavy overfit. The method is based on a new empirical result, which shows that the agreement among deep networks decreases with the occurrence of overfit. Further support is provided by the theoretical analysis of a linear regression model.

Our empirical study focused for the most part, though not entirely, on datasets with noisy labels, a real-world condition that creates heavy overfitting, where our method is most useful. In almost all cases, the method’s resistance to overfit is maintained even when the level of noise is very high. In such cases, our method significantly outperforms regular ensembles, as well as other post-processing methods designed specifically to handle noisy labels. Its advantage over early stopping, shown in Fig. 5, is particularly interesting, as it shows that our method could, in some cases, allow the user to harness the new “knowledge” learned by the model at the late stages of training - even after overfit has occurred - while minimizing the damages of overfit.

Our method has some practical advantages. Notably, it is easy to implement and can be readily used at post-processing with almost any method, network, and dataset. This is accomplished *without any change to the training process and without any additional hyper-parameter tuning*. Since our method is only employed at post-processing, when no overfit is suspected it can be completely avoided. Thus, it serves as a practical tool to combat overfit (with or without label noise in the training set), especially when large amounts of correctly labeled data are difficult to acquire.

## References

- Arpit, D.; Jastrzebski, 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*, 233–242. PMLR.
- Bae, H.; Shin, S.; Na, B.; Jang, J.; Song, K.; and Moon, I.-C. 2022. From Noisy Prediction to True Label: Noisy Prediction Calibration via Generative Model. In *International Conference on Machine Learning*, 1277–1297. PMLR.
- Breiman, L. 1996. Bagging predictors. *Machine learning*, 24(2): 123–140.
- Freund, Y.; and Schapire, R. E. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1): 119–139.
- Ganaie, M. A.; Hu, M.; et al. 2021. Ensemble deep learning: A review. *arXiv preprint arXiv:2104.02395*.
- Hacohen, G.; Choshen, L.; and Weinshall, D. 2020. Let's Agree to Agree: Neural Networks Share Classification Order on Real Datasets. In *Int. Conf. Machine Learning ICML*, 3950–3960.
- Hansen, L. K.; and Salamon, P. 1990. Neural network ensembles. *IEEE transactions on pattern analysis and machine intelligence*, 12(10): 993–1001.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456. PMLR.
- Karimi, D.; Dou, H.; Warfield, S. K.; and Gholipour, A. 2020. Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis. *Medical Image Analysis*, 65: 101759.
- Krogh, A.; and Hertz, J. 1991. A simple weight decay can improve generalization. *Advances in neural information processing systems*, 4.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Lee, K.; Yun, S.; Lee, K.; Lee, H.; Li, B.; and Shin, J. 2019. Robust inference via generative classifiers for handling noisy labels. In *International Conference on Machine Learning*, 3763–3772. PMLR.
- Li, H.; Wang, X.; and Ding, S. 2018. Research and development of neural network ensembles: a survey. *Artificial Intelligence Review*, 49(4): 455–479.
- Liu, S.; Niles-Weed, J.; Razavian, N.; and Fernandez-Granda, C. 2020. Early-learning regularization prevents memorization of noisy labels. *Advances in neural information processing systems*, 33: 20331–20342.
- Nakkiran, P.; Kaplun, G.; Bansal, Y.; Yang, T.; Barak, B.; and Sutskever, I. 2021. Deep double descent: Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12): 124003.
- Nicholson, B.; Sheng, V. S.; and Zhang, J. 2016. Label noise correction and application in crowdsourcing. *Expert Systems with Applications*, 66: 149–162.
- Patrini, G.; Rozza, A.; Krishna Menon, A.; Nock, R.; and Qu, L. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1944–1952.
- Pliushch, I.; Mundt, M.; Lupp, N.; and Ramesh, V. 2021. When Deep Classifiers Agree: Analyzing Correlations between Learning Order and Image Statistics. *arXiv preprint arXiv:2105.08997*.
- Salman, S.; and Liu, X. 2019. Overfitting mechanism and avoidance in deep neural networks. *arXiv preprint arXiv:1901.06566*.
- Shorten, C.; and Khoshgoftaar, T. M. 2019. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1): 1–48.
- Shwartz, D.; Stern, U.; and Weinshall, D. 2022. The Dynamic of Consensus in Deep Networks and the Identification of Noisy Labels. *arXiv preprint arXiv:2210.00583*.
- Song, H.; Kim, M.; Park, D.; Shin, Y.; and Lee, J.-G. 2022. Learning from noisy labels with deep neural networks: A survey. *IEEE Transactions on Neural Networks and Learning Systems*.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958.
- Stephenson, C.; and Lee, T. 2021. When and how epochwise double descent happens. *arXiv preprint arXiv:2108.12006*.
- Stern, U.; Shwartz, D.; and Weinshall, D. 2023. United We Stand: Using Epoch-wise Agreement of Ensembles to Combat Overfit. *arXiv preprint arXiv:2310.11077*.
- Stern, U.; and Weinshall, D. 2023. Relearning Forgotten Knowledge: on Forgetting, Overfit and Training-Free Ensembles of DNNs. *arXiv preprint arXiv:2310.11094*.
- Wen, Y.; Tran, D.; and Ba, J. 2020. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. *arXiv preprint arXiv:2002.06715*.
- Wenzel, F.; Snoek, J.; Tran, D.; and Jenatton, R. 2020. Hyperparameter ensembles for robustness and uncertainty quantification. *Advances in Neural Information Processing Systems*, 33: 6514–6527.
- Zhang, M.; Lee, J.; and Agarwal, S. 2021. Learning from noisy labels with no change to the training process. In *International Conference on Machine Learning*, 12468–12478. PMLR.