
Generalized Negative Correlation Learning for Deep Ensembling

Sebastian Buschjäger¹ Lukas Pfahler¹ Katharina Morik¹

Abstract

Ensemble algorithms offer state of the art performance in many machine learning applications. A common explanation for their excellent performance is due to the bias-variance decomposition of the mean squared error which shows that the algorithm’s error can be decomposed into its bias and variance. Both quantities are often opposed to each other and ensembles offer an effective way to manage them as they reduce the variance through a diverse set of base learners while keeping the bias low at the same time. Even though there have been numerous works on decomposing other loss functions, the exact mathematical connection is rarely exploited explicitly for ensembling, but merely used as a guiding principle. In this paper, we formulate a generalized bias-variance decomposition for arbitrary twice differentiable loss functions and study it in the context of Deep Learning. We use this decomposition to derive a Generalized Negative Correlation Learning (GNCL) algorithm which offers explicit control over the ensemble’s diversity and smoothly interpolates between the two extremes of independent training and the joint training of the ensemble. We show how GNCL encapsulates many previous works and discuss under which circumstances training of an ensemble of Neural Networks might fail and what ensembling method should be favored depending on the choice of the individual networks.

1. Introduction

Ensemble algorithms offer state of the art performance in many Machine Learning applications and often outperform single classifiers by a large margin. One of the main theoretical driving forces behind the understanding of ensembles is the bias-variance decomposition. The bias-variance decom-

position decomposes the algorithm’s error into two additive parts – its bias and its variance. Hence, a good algorithm should try to minimize both at the same time which often leads to a difficult balancing act between the two quantities. Ensemble algorithms are well-known to reduce the variance if a diverse set of base models is trained while also keeping the bias low. The bias-variance decomposition has been mathematically proven for the mean-squared error, which sparked a plethora of different learning for different loss functions exploiting the general notion of ‘diversity’ in ensemble construction (Webb, 2000; Geurts et al., 2006; Brown et al., 2005; Melville & Mooney, 2005; Lee et al., 2015; Zhou & Feng, 2017; Dvornik et al., 2019). Interestingly, even though there have been numerous works on decomposing other loss functions such as 0 – 1 loss or exponential families (e.g. log-likelihood loss), none of these theoretical insights have directly inspired new learning algorithms. The general notion of diversity in an ensemble is still one of the main driving forces in designing new ensembling algorithms, while the exact mathematical connection is rarely exploited explicitly.

We argue, that diversity can be hurtful sometimes and must be controlled with respect to the base learners. To do so, we formulate a generalized bias-variance decomposition and study it in the context of Deep Learning. From this decomposition, we derive two different algorithmic extremes: Either, we train the entire ensemble jointly in an end-to-end fashion or we train each model completely independent from each other. We present a generalization of Negative Correlation Learning (GNCL) that smoothly interpolates between these two extremes and thus can capitalize on the entire spectrum of methods in-between. We show, how GNCL generalizes many existing ensembling techniques in a single framework and use it to explore under which circumstances training of an ensemble might fail and what ensembling method should be favored depending on the choice of the individual networks. Our contributions are:

- **A Generalized Bias-Variance Decomposition:** We present the first bias-variance decomposition for arbitrary twice differentiable loss functions.
- **Generalized Negative Correlation Learning:** From this decomposition we derive Generalized Negative Correlation Learning (GNCL) and show how it gener-

¹Artificial Intelligence Group, TU Dortmund University, Germany. Correspondence to: Sebastian Buschjäger <sebastian.buschjaeger@tu-dortmund.de>.

alizes existing NCL-like algorithms into a single framework.

- **Experimental evaluation:** We compare our approach against state-of-the-art ensemble algorithms for Deep Learning methods. We show how GNCL smoothly interpolates between different ensembling techniques offering the overall best performance.
- **Explanation of results:** Our theoretical results accurately explain when certain ensembling methods should be favored over others: For small capacity Neural Networks, End-to-End learning should be favored, whereas, for larger capacity models, ensembling should shift towards independent training.

The paper is organized as follows: The next section surveys related work and focuses on the bias-variance decomposition as well as ensembling methods in the realm of Deep Learning. Section 3 then derives the bias-variance decomposing, whereas section 4 formalizes it into the Generalized Negative Correlation Learning algorithm. In section 5 we experimentally evaluate our method and section 6 concludes the paper.

2. Related Work

The first bias-variance decomposition was proposed by Harry Markowitz in (Markowitz, 1952) for the mean squared error (MSE), which was later found to be one of the cornerstones of modern financial portfolio theory. Its first appearance - also for the MSE - in the Machine Learning community was due to Geman et al. in (Geman et al., 1992) which then sparked a series of different decompositions (see e.g. (Ueda & Nakano, 1996; Domingos, 2000; James, 2003) and references therein). Most notable is the work by Domingos in (Domingos, 2000) as it provides a set of consistent definitions for bias and variance and gives rise to a decomposition of the 0 – 1 loss which fits the previous decompositions of the MSE. However, we note that these decompositions either focus on the mean squared error or the 0 – 1 loss for binary classification problems, but not for general loss functions. James claims in (James, 2003) to give a generalized bias-variance decomposition for any symmetric loss functions by providing a set of different definitions of Bias and Variance. He then continues to choose those definitions which ‘fit’ the original MSE decomposition best but never proves the consistency of these definitions. Moreover, as he notes, these definitions are only applicable for binary classification problems and not applicable for real-valued predictions. A similar decomposition has also been proposed in the context of Product Of Expert ensembles called the ambiguity decomposition. This decomposition also first appeared for the MSE and is equal to the bias-variance decomposition although derived from a distributional point

of view (Krogh & Vedelsby, 1995; Heskes, 1998). Later, Hansen and Heskes give in (Hansen & Heskes, 2000) a generalized ambiguity decomposition for exponential families. The authors assume that an additive decomposition with two summands of the form Error = Bias + Variance exist and then continue to show that an exponential family will always result in such a decomposition. They do not discuss distributions which do not fit this assumption. Most closely related to our approach is the work due to Jiang et al. in (Jiang et al., 2017). Here, the authors derive a Generalized Ambiguity Decomposition for twice differentiable loss functions. Similar to our approach, the authors also use a second-order Taylor approximation around the ensemble’s prediction $f(x)$ but seem to ignore the remainder in their construction. This paper focuses on binary classification losses with a single output and does not directly translate into a new learning algorithm. Our approach on the other hand also encapsulates multi-class problems and therefore is a natural generalization of previous work. Moreover, we present a novel learning objective and show how this objective encapsulates many well-known existing objectives presented in the literature.

Even though the exact theoretical connection between Bias and Variance for other loss functions was missing, the decomposition for the MSE sparked a multitude of different algorithms. In this paper, we focus on ensembles of Deep Nets and Neural Networks. However, we note that our generalized bias-variance decomposition does not assume any specific base learners, but is equally applicable to any base learner, e.g. Decision Trees. In the realm of Neural Networks, Negative Correlation Learning (NCL) is a direct application of the bias-variance decomposition and was first proposed by Liu et al. in (Liu & Yao, 1999) and later refined by Brown et al. in (Brown et al., 2005). Opitz et al. use NCL as inspiration to enforce diversity among neural networks in an ensemble by employing the cross-entropy loss between the individual experts’ outputs (Opitz et al., 2016). Dvornik et al. propose in (Dvornik et al., 2019) a similar ensembling technique but are more freely in their choices to enforce diversity. Specifically, they train each network on the cross-entropy loss but employ the cosine-similarity and KL-Divergence as a regularization term to enforce diversity. Webb et al. recently proposed a similar objective which they justify by viewing the ensemble as a product of experts. This leads to the minimization of the KL-Divergence to preserve most of the contribution of each expert with a coupling term similar to NCL (Webb et al., 2019; 2020). Lee et al. train in (Lee et al., 2015; 2016) a diverse ensemble of classifiers by using Stochastic Multiple Choice Learning (SMCL). Instead of training all ensemble members on all the available data, they only update that member with the smallest loss. This way, the diversity which naturally occurs due to the random initialization, is promoted. Bagging

(Breiman, 1996) has also been applied in the context of Deep Learning (Brown et al., 2005; Zeiler & Fergus, 2014; Lee et al., 2015; Lakshminarayanan et al., 2017; Zhu et al., 2019; Ovadia et al., 2019; Webb et al., 2019; 2020) which enforces diversity by training each expert individually on a bootstrap sample of the training data¹. Some works argue, that the random initialization of Deep Nets combined with stochastic gradient descent promotes enough diversity (Lee et al., 2015; Lakshminarayanan et al., 2017; Ovadia et al., 2019; Devlin et al., 2019) so that bootstrap samples are not required. It is also noteworthy, that this training method sometimes occurs as a special case for certain hyperparameter settings (Brown et al., 2005; Webb et al., 2019; 2020), including this work. Joint training of the entire ensemble in an End-to-End fashion has also been proposed (Brown et al., 2005; Webb et al., 2019; 2020; Opitz et al., 2016; Dutt et al., 2017; Lee et al., 2015). This approach ignores the bias and variance of the individual experts but focuses on the ensemble’s joint loss. Here, the literature is slightly more fragmented. End-To-End training also occurs in (Brown et al., 2005; Opitz et al., 2016; Webb et al., 2019; 2020) as a special case for certain hyperparameter settings, including this work. Dutt et al. call this approach a ‘coupled ensemble’ (Dutt et al., 2017), whereas Lee et al. call this approach training under an ‘ensemble-aware’ loss (Lee et al., 2015). While less extensive, Boosting (Schapire & Freund, 2012) has also been applied to Neural Networks and Deep Learning. Early works focused on the combination of smaller Neural Networks as base learners for ensembling which also carried over to larger architectures commonly found in Deep Learning (Opitz & Maclin, 1997; Maclin & Opitz, 1997; Schwenk & Bengio, 2000; Moghimi et al., 2016; Zhu et al., 2019). For reference, we note that there has also been some interest in understanding Residual architectures (ResNet) as boosting in feature space (Huang et al., 2018). Recently, ensembles which are derived from a single network have also been proposed. Dropout (Srivastava et al., 2014) is anecdotally sometimes referred to as ‘the ensemble of possible subnetworks’ (Baldi & Sadowski, 2013; Gal & Ghahramani, 2016). This connection has been studied more closely in the context of ‘pseudo-ensembles’ (Bachman et al., 2014). Pseudo-ensembles are ensembles that are derived from a large single network by perturbing it with a noise process, e.g. by removing weights as done by Dropout. Although not explicitly mentioned, ‘snapshot ensembles’ (Qiu et al., 2014; Huang et al., 2017a) which store multiple versions of the same network during the optimization (e.g. by storing the current model every 10 epochs) can also be seen in this framework.

¹To the best of our knowledge, there is no publication on feature bagging for training Deep Networks.

3. A Generalized Bias-Variance Decomposition

We consider a supervised learning setting, in which we assume that training and test points are drawn i.i.d. according to some distribution \mathcal{D} over the input space \mathcal{X} and labels \mathcal{Y} . For training, we have given a labelled sample $\mathcal{S} = \{(x_i, y_i) | i = 1, \dots, N\}$, where $x_i \in \mathcal{X} \subseteq \mathbb{R}^d$ is a d -dimensional feature-vector and $y_i \in \mathcal{Y} \subseteq \mathbb{R}^C$ is the corresponding target vector. For binary classification problems we set $C = 1$ and $\mathcal{Y} = \{-1, +1\}$; for regression problems we have $C = 1$ and $\mathcal{Y} = \mathbb{R}$. For multiclass problems with C classes we encode each label as a one-hot vector $y = (0, \dots, 0, 1, 0, \dots, 0)$ which contains a ‘1’ at coordinate c for label $c \in \{0, \dots, C - 1\}$.

Given a model class $\mathcal{H} = \{h: \mathcal{X} \rightarrow \mathbb{R}^C\}$ we wish to select that model that fits our current sample. In practice we employ various learning algorithms to do so, e.g. SGD in the context of Deep Learning or CART in the context of decision trees. These algorithms often introduce some form of randomization, e.g. by random initialization of weights or random sampling of splits thus introducing some model distribution Θ . Yet, in the heart of these algorithms, we find the minimization of a loss function $\ell: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ which quantifies the error of our model’s prediction $h(x)$ compared to the real target y . So to choose the optimal algorithm for a problem, we may favor that algorithm which consistently produces the best models with the smallest loss:

$$\Theta^* = \arg \min_{\Theta} \mathbb{E}_{h \sim \Theta, (x, y) \sim \mathcal{D}} [\ell(h(x), y)]$$

For the rest of this paper we assume that ℓ is at least twice continuous differentiable and present appropriate choices at the end of this section. We now use a second-order Taylor approximation of ℓ around the centre $\mu(x) = \mathbb{E}_{h \sim \Theta} [h(x)]$. For readability we now drop the subscript $h \sim \Theta, (x, y) \sim \mathcal{D}$. Similarly we write $h(x) = h$ and $\mu(x) = \mu$ and $\ell(h(x), y) = \ell(h)$:

$$\begin{aligned} \mathbb{E} [\ell(h)] &= \mathbb{E} [\ell(\mu)] + \mathbb{E} [(h - \mu)^T \nabla_{\mu} \ell(\mu)] \\ &\quad + \mathbb{E} \left[\frac{1}{2} (h - \mu)^T \nabla_{\mu}^2 \ell(\mu) (h - \mu) \right] + \mathbb{E} [R_3] \end{aligned}$$

where R_3 denotes the remainder of the Taylor approximation containing the third and higher derivatives.

We note, that $\nabla_{\mu} \ell(\mu)$ does not depend on h since μ is a constant given a fixed test point x and therefore $\mathbb{E} [\nabla_{\mu} \ell(\mu)] = \nabla_{\mu} \ell(\mu)$. Also note, that per definition $\mathbb{E} [h] = \mu$ so that the second summand vanishes:

$$\mathbb{E}_h [(h - \mu)^T \nabla_{\mu} \ell(\mu)] = 0$$

Naturally, the quality of this approximation depends on the magnitude of the remainder and it becomes exact if the loss function does not have a third derivative. Otherwise, we

may use a classic text-book results to bound the magnitude of the remainder for functions which are 3 times continuous differentiable. Let there be some $m \in \mathbb{R}$ so the third derivative of the loss is bounded by it, that is $|\nabla_\mu^3 \ell(\mu)|_{i,j,k} \leq m$ for all $(x, y) \sim \mathcal{D}$ then

$$R_3(h-\mu) \leq \frac{1}{6}m \max_h \|h-\mu\|_1^3 \leq \frac{1}{6}mC \max_{h_1, \dots, h_C} (h_i - \mu_i)^3$$

For a sufficiently small remainder we approximate:

$$\mathbb{E} [\ell(h)] \approx \mathbb{E} [\ell(\mu)] + \mathbb{E} \left[\frac{1}{2} \phi^T \nabla_\mu^2 \ell(\mu) \phi \right] \quad (1)$$

$$= \mathbb{E} [\ell(\mu)] + \frac{1}{2} \text{tr} (\nabla_\mu \ell(\mu) \text{cov}(\phi, \phi)) \quad (2)$$

where $\phi = (h - \mu)$ and the second line is the quadratic form of the expectation. We may interpret this decomposition as a generalized Bias-(Co-)Variance decomposition: While the LHS depicts the expected error of a model h , the first term on the RHS depicts the error of the expected model - or differently coined the *algorithm's* bias. Moreover, the second term can be interpreted as the co-variance of h with respect to the expected model μ given a loss-specific multiplicative constant $\nabla^2 \ell(\mu)$.

3.1. Example 1: Mean-squared error

Consider the mean squared error (MSE) of a one dimensional regression task $\mathcal{Y} = \mathbb{R}$ and let $z = h(x)$:

$$\begin{aligned} \ell(z, y) &= \frac{1}{2}(z - y)^2 \\ \frac{\partial \ell}{\partial z} &= (z - y) \\ \frac{\partial^2 \ell}{\partial z \partial h(x)} &= 1 \\ \frac{\partial^3 \ell}{\partial z \partial h(x) \partial z} &= 0 \end{aligned}$$

The third derivative of the MSE vanishes and thus the above approximation is exact. The resulting decomposition matches exactly the well-known Bias-Co-Variance decomposition.

3.2. Example 2: Negative-likelihood Loss

As a third example we consider multi-class classification problem with C classes. Let $z = h(x) \in \mathbb{R}^C$ and let ℓ be

the negative-likelihood loss (NLL):

$$\begin{aligned} \ell(z, y) &= - \sum_{i=1}^C y_i \log(z_i) \\ \frac{\partial \ell}{\partial z_i} &= - \frac{y_i}{z_i} \\ \frac{\partial^2 \ell}{\partial z_i \partial z_j} &= \frac{y_i}{z_i^2} \mathbb{1}\{i = j\} \\ \frac{\partial^3 \ell}{\partial z_i \partial z_j \partial z_k} &= -2 \frac{y_i}{z_i^3} \mathbb{1}\{i = j = k\} \end{aligned}$$

For this loss function, the third derivative does not vanish and thus the decomposition is not exact. Looking at the third derivative we also see, that it can get uncontrollably large for $z_i \rightarrow 0$ if $y_i = 1$. Thus, if a model completely fails with a wrong prediction then the decomposition error can be unbounded. Put differently, the performance of a model using the NLLLoss cannot be completely explained in terms of ‘Bias’ and ‘Variance’ since the remainder is not neglectable.

3.3. Example 3: Cross Entropy Loss

A common combination in Deep Learning is the NLLLoss with the softmax function. Again, let $z = h(x) \in \mathbb{R}^C$. The softmax function maps each output dimension z_i of the classifier to a probability:

$$q_i = \frac{e^{z_i}}{\sum_{i=1}^C e^{z_j}}$$

We combine softmax with the NLLLoss:

$$\begin{aligned} \ell(z, y) &= - \sum_{i=1}^C y_i \log(z_i) = - \sum_{i=1}^C y_i \log \left(\frac{e^{z_i}}{\sum_{i=1}^C e^{z_j}} \right) \\ \frac{\partial \ell}{\partial z_i} &= q_i - \mathbb{1}\{y_i = 1\} \\ \frac{\partial^2 \ell}{\partial z_i \partial z_j} &= q_i (\mathbb{1}\{i = j\} - q_j) \\ \frac{\partial^3 \ell}{\partial z_i \partial z_j \partial z_k} &= \mathbb{1}\{i = j\} q_i (\mathbb{1}\{i = k\} - q_k) \\ &\quad - q_i q_j (\mathbb{1}\{i = k\} - q_k) \\ &\quad - q_i q_j (\mathbb{1}\{j = k\} - q_k) \end{aligned}$$

Due to the softmax function we have $\sum_{c=1}^C q_c = 1, q_c > 0 \forall c = 1, \dots, C$. The maximum of the third derivative is obtained for pairwise unequal i, j, k ($i \neq j, j \neq k, i \neq k$) and $q_i = q_j = q_k = \frac{1}{3}$:

$$2 \cdot q_i q_j q_k \leq \frac{1}{27} < 0.038$$

Thus, the decomposition error for the cross entropy loss is bounded and we can explain a models performance in terms of its Bias and Variance (up to the bounded remainder).

4. Generalized Negative Correlation Learning

As often faced in Machine Learning we cannot compute $\mathbb{E}_{x,y \sim \mathcal{D}} [\ell(h(x), y)]$ exactly since we do not know the exact distribution \mathcal{D} and in fact, this is part of the problem we would like to solve. Moreover, it is difficult to compute $\mu(x) = \mathbb{E}_{h \sim \Theta} [h(x)]$ exactly since the algorithm we use for computing h (e.g. SGD) only implicitly induces a distribution over h and the exact nature of Θ for various model classes is ongoing research (Biau & Scornet, 2015; Sutskever et al., 2013; Arora et al., 2019; Kawaguchi et al., 2017). For sufficiently large training sample \mathcal{S} we use Monte-Carlo approximation:

$$\mathbb{E}_{x,y} [\ell(h(x), y)] \approx \frac{1}{N} \sum_{(x,y) \in \mathcal{S}} \ell(h(x), y)$$

Similarly, we may approximate the expected prediction μ with M models:

$$\begin{aligned} \mu(x) &= \mathbb{E}_{\Theta} [h(x)] \approx f(x) = \frac{1}{M} \sum_{i=1}^M h^i(x) \\ \mathbb{E}_{\Theta} \left[\frac{1}{2} \phi^T \nabla_{\mu}^2 \ell(\mu) \phi \right] &\approx \frac{1}{2M} \sum_{i=1}^M d_i^T D d_i \\ \mathbb{E}_{\Theta} [R(x)] &\approx \tilde{R} \end{aligned}$$

where $D = \nabla_{f(x)}^2 \ell(f(x), y)$ and $d_i = (h^i(x) - f(x))$. We stress the fact, that we *assume* that these are good approximations. For large M , this is certainly a justified approximation, but for smaller M this is not necessarily the case. However, additive ensembles of this form are arguably the most common form of ensembles and undeniably work well in practice. We define the empirical bias-variance decomposition for any twice-differentiable loss function as:

$$\ell(f) = \frac{1}{M} \sum_{i=1}^M \ell(h^i) - \frac{1}{2M} \sum_{i=1}^M d_i^T D d_i + \tilde{R} \quad (3)$$

We note, that for any convex loss function D is positive definite and therefore $d_i^T D d_i \geq 0$ which implies:

$$\ell(f) \leq \frac{1}{M} \sum_{i=1}^M \ell(h^i)$$

It follows, that an ensemble of models will always be better than a single model making a compelling argument for ensemble learning. Note, that a similar argument has been made numerous times already and can for example directly be obtained when applying Jensen's inequality to the weighted average of models over a convex loss.

We use Eq. 3 as a basis for a learning algorithm: We can either directly minimize its LHS and optimize the entire ensemble in an end-to-end fashion. Alternatively, we use

its RHS to derive a regularized objective which trains each network independently with a coupling term enforcing some diversity. To do so, let \tilde{R} be sufficiently small, so that we may ignore it and let $\lambda \in \mathbb{R}$ be a regularization parameter, then we may minimize:

$$\frac{1}{M} \sum_{i=1}^M \ell(h^i) - \frac{\lambda}{2M} \sum_{i=1}^M d_i^T D d_i \quad (4)$$

4.1. A combined loss function

Having the two objectives $\ell(f)$ and eq. 4 available begs the question of which one may lead to better results. Frankly, since both objectives are equal, minimizing both will lead to similar if not equal results. Thus, using either approach comes down to the more practical specifics of the problem at hand: Direct minimization of the loss seems favorable because it automatically finds a good trade-off between bias and variance and no hyperparameter tuning is necessary. Yet, using Eq. 4 on the other hand enables us to train each model independently and only requires some synchronization between models to make sure that the variance is large enough (See e.g. (Webb et al., 2019) and references therein for a discussion on distributed training). Moreover, this approach allows practitioners to fine-tune the trade-off between bias and variance which might be favorable for specific problems and base models. In Deep Learning, it is common practice to train networks to achieve zero loss on the training data and sometimes train it even longer (Zhang et al., 2017). Recall that for a convex loss $d_i^T D d_i \geq 0$ and therefore Eq. 3 implies that (for a sufficiently small remainder) an ensemble with powerful base learners having zero training loss should not have any variance on the training data. Therefore, as soon as the base learners achieve zero training loss there is no need to invest into variance because the best model (from the training data's perspective) has already been found. Clearly, this is neither the intuition behind the bias-variance decomposition nor is it what we are trying to achieve. And indeed, in most practical applications we can be sure that even though we have zero training loss, that we will suffer some loss when applying our model to new, unseen data. In this case, it might still be favorable to enforce some diversity between base models during training to achieve a better generalization error. We will investigate this effect in our experiments more carefully and show that there is a clear dependence on which method to favor depending on the type and strength of the base learner.

Interestingly, there is an upper bound of the bias-variance decomposition that combines both approaches into a single objective. This upper bound simply re-scales the individual contributions of the base learners and thus results in the *same* solution as minimizing $\ell(f)$ or Eq. 4 for appropriate choices of λ which allows us to smoothly interpolate between the two extremes of independent and end-to-end

training. In addition, this formulation circumvents the costly computation of D and does not require the assumption that \tilde{R} is sufficiently small:

$$\begin{aligned} & \frac{1}{M} \sum_{i=1}^M \ell(h^i) - \frac{1}{2M} \sum_{i=1}^M d_i^T D d_i + \tilde{R} \\ & \leq \frac{1}{M} \sum_{i=1}^M \ell(h^i) - \frac{1}{2M} \sum_{i=1}^M d_i^T D d_i + \tilde{R} + \frac{1}{M} \sum_{i=1}^M \ell(h^i) \\ & = \ell(f) + \frac{1}{M} \sum_{i=1}^M \ell(h^i) \end{aligned}$$

To this end we propose the following Generalized Negative Correlation Learning (GNCL) objective for $\lambda \in [0, 1]$:

$$\frac{1}{N} \sum_{j=1}^N \left(\lambda \ell(f(x_j), y_j) + \frac{1-\lambda}{M} \sum_{i=1}^M \ell(h^i(x_j), y_j) \right) \quad (5)$$

For $\lambda = 0$ this trains M models independently, whereas for $\lambda = 1$ all models are trained jointly in an end-to-end fashion. For values between zero and one we can smoothly interpolate between these to extremes making the entire spectrum of trade-offs available.

4.2. Relationship to other ensembling approaches

There are multiple mentions of NCL-like algorithms in literature. We will now show, that these algorithms are a special version of the proposed Generalized Negative Correlation Learning algorithm.

Negative Correlation Learning: The earliest works (Liu & Yao, 1999; Brown et al., 2005) on NCL-Learning propose to minimize the MSE with a coupling term including the ensembles' diversity (c.f. Eq. (17) in (Brown et al., 2005)):

$$\frac{1}{M} \sum_{i=1}^M \frac{1}{2} (h^i(x) - y)^2 - \lambda \frac{1}{M} \sum_{i=1}^M \frac{1}{2} (h^i(x) - f(x))^2$$

Substituting the second derivative of the MSE loss in Eq. 4 directly leads to this formulation. NCL is a specialized version of GNCL for the MSE loss.

Modular loss: Webb et al. propose to minimize both, the ensemble loss as well as the loss of each individual expert in a modular loss function (c.f. Eq (4) in (Webb et al., 2019)):

$$\lambda D(f(x)||y) + (1-\lambda) \frac{1}{M} \sum_{i=1}^M D(h^i(x)||y)$$

where D denotes the KL-Divergence and $\lambda \in [0, 1]$ is the regularization strength. Substituting the cross-entropy loss into Eq. 5 yields the same formulation. The modular loss is a specialized version of GNCL with the cross entropy loss.

DivLoss: Opitz et al. use NCL as inspiration to enforce diversity among neural networks by employing the cross-entropy loss between the individual experts' outputs while minimizing the individual and the ensemble loss. They propose to minimize (c.f. Eq. (15) in (Opitz et al., 2016)) the DivLoss:

$$\begin{aligned} & \ell(f(x), y) + \frac{\lambda_1}{M} \sum_{i=1}^M \ell(h^i(x), y) \\ & \quad - \frac{\lambda_2}{M(M-1)} \sum_{i=1}^M \sum_{j \neq i} \ell(h^i(x), h^j(x)) \end{aligned}$$

where ℓ is the cross-entropy loss with softmax activation and $\lambda_1, \lambda_2 \in \mathbb{R}$ are regularization parameters. Note that $\ell(f(x), y) \geq 0$ and that $\ell(h^i(x), h^j(x))$ is convex in its first argument. For a fixed scaling $\kappa \leq 1$ it holds that

$$\kappa \frac{1}{M} \sum_{i=1}^M \ell(h^i(x), y) \leq \ell \left(\frac{1}{M} \sum_{i=1}^M h^i(x), y \right)$$

due to Jensen's inequality. Therefore, we lower-bound the original objective for $\lambda_2 \leq 1$ to:

$$\frac{\lambda_1}{M} \sum_{i=1}^M \ell(h^i(x), y) - \frac{\lambda_2}{(M-1)} \sum_{j=1}^M \ell(f(x), h^j(x))$$

Interestingly, Webb et al. show in (Webb et al., 2019) that this formula is an alternative formulation of their modular loss when setting $\lambda_1 = 1$ and $\lambda_2 = \lambda \in [0, 1]$. It follows, the objective proposed in (Opitz et al., 2016) is an upper bound of the modular loss proposed in (Webb et al., 2019), which in turn is a specialized version of GNCL learning for the cross entropy loss.

Diversity with Cooperation: Dvornik et al. propose in (Dvornik et al., 2019) an ensemble approach that focuses on diversity and cooperation at the same time. More formally, they propose to use the following objective

$$\sum_{i=1}^M \ell(h^i(x), y) + \frac{\lambda}{(M-1)} \sum_{i=1}^M \sum_{j \neq i} \psi(h^i(x), h^j(x))$$

where ψ is a penalty function to enforce diversity in the ensemble. By using the cross entropy loss and setting $\psi = -\ell$ we arrive at the DivLoss function for $\lambda_1 = M$ and $\lambda_2 = 1$. Thus, the diversity with cooperation approach by Dvornik et al. is closely related to GNCL. However, we note that the authors are freer with their choices of ψ leading to mixed experimental results.

Bagging and Wagging: Bagging uses bootstrap samples to assign a different subset of training examples to each expert and thereby enforces diversity. Bauer and Kohavi

(Bauer & Kohavi, 1999) propose an extension called Wagging which samples different weights instead of sampling examples directly. Oza and Russel show in (Oza & Russell, 2001) that Wagging with weights sampled from a discrete Poisson distribution $w \sim \text{Poisson}(1)$ is the same as Bagging. Similarly, Webb et al. propose in (Webb, 2000) to use continuous Poisson weights for Wagging, which improves the performance for certain base learners. Formally, the loss function for Wagging and its variants is

$$\frac{1}{N} \sum_{j=1}^N \frac{1}{M} \sum_{i=1}^M w_{j,i} \tilde{\ell}(h^i(x_j), y_j)$$

where $w_{j,i} \in \mathbb{R}_+$ is the precomputed weight for each base learner and sample and $\tilde{\ell}: \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ is another loss function. Setting $\ell(h^i(x_j), y_j) = w_{i,j} \tilde{\ell}(h^i(x_j), y_j)$ and $\lambda = 0$ in Eq. 5 yields the same formulation. Hence, Wagging and Bagging are specialized versions of GNCL for weighted loss functions.

Boosting: Boosting iteratively trains new classifiers to correct the errors of the previous classifier, thereby constructing a strong classifier from weak base models. It is well-known that Boosting can be viewed as functional gradient descent in which each new base learner tries to approximate the negative gradient of a loss function (Mason et al., 2000; Schapire & Freund, 2012). Thus, boosting minimizes the entire ensemble loss and is therefore closely related to GNCL with $\lambda = 1$. However, we note that Boosting fundamentally behaves differently from GNCL because it is designed to greedily approximate gradients for non-differential base learner’s functions. Theoretically, both approaches could be combined: The proposed GNCL objective can either be minimized via Stochastic Gradient Descent (as done in our experiments) or by boosting weak-learner on the GNCL objective. Last, note that Residual Networks (ResNet) have also been shown to share a connection with Boosting. ResNets perform boosting in feature space with a telescoping sum instead of a weighted average as used by Boosting (Huang et al., 2018) and thus – while there is some overlap – both methods are fundamentally different. Interestingly, we found that ResNet architectures performed best as a base learner during our experiments.

Dropout and Pseudo-Ensembles: Dropout (Srivastava et al., 2014) is a regularization method for Deep Nets, which randomly sets weights to zero during the forward pass. While Dropout helps to prevent overfitting, it can also be used to estimate the geometric mean and variance of a distribution of networks with parameter sharing. This is anecdotally sometimes referred to as ‘the ensemble of possible subnetworks’ (Baldi & Sadowski, 2013; Gal & Ghahramani, 2016). Bachman et al. studied this connection more closely and proposed in (Bachman et al., 2014) the term ‘pseudo-ensembles’. Pseudo-ensembles are ensembles that

are derived from a large single network by perturbing it with a noise process, e.g. by removing weights as done by Dropout. Although not explicitly mentioned, snapshot ensembles (Qiu et al., 2014; Huang et al., 2017a) which store multiple versions of the same network (e.g. by storing the current model every 10 epochs) can also be seen in this framework. Pseudo-ensembles minimize the following objective

$$\frac{1}{N} \sum_{j=1}^N \mathbb{E}_{\theta} [\ell_{\theta}(\mu(x_i), y_i)] + \lambda \mathbb{E}_{\theta} [R(\mu(x_i), \mu_{\theta}(x_i))]$$

where μ denotes the ‘mother’ net, μ_{θ} is a child net under the noise process θ , ℓ is a loss function and R is a regularizer with regularization strength λ . Note, that for our bias-variance decomposition we derived the same objective in eq. 1 with $R = \phi^T \nabla_{\mu}^2 \ell(\mu) \phi$ and by introducing λ as discussed earlier. Unfortunately, the authors do not discuss how to directly minimize this objective under the noise process θ . Interestingly, for experiments they use the *same* formulation as our GNCL objective in Eq. 5 with the cross entropy loss². We conclude, that the proposed bias-variance decomposition also encapsulates pseudo-ensembles and the GNCL objective can be viewed as an empirical version of this. However, we note that Pseudo-Ensembles have a very different viewpoint to our approach: Pseudo-Ensembles train a single network and spawns a diverse set of offsprings from this large network, whereas GNCL combines a set of smaller models into a large one.

5. Experiments

In our experimental evaluation, we study two different aspects of the generalized bias-variance decomposition. As discussed before, when the base learners achieve zero training loss then the bias-variance decomposition implies, that the ensemble should not have any diversity. However, when applied to new, unseen data the base learners will likely have a nonzero loss and therefore, for a better generalization, it might be worthwhile to enforce some diversity during training. We will study this phenomenon in the context of Deep Learning by training ensembles with base learners of different capacities. For our evaluation, we use the CIFAR100 dataset (Krizhevsky, 2009) which contains 50000 $32 \times 32 \times 3$ images of various everyday objects which belong to one of 100 classes. For testing, we utilize the given test split with 10000 images. In all experiments, we perform standard data augmentation during training (Random cropping, random horizontal flipping, and normalization). We train for 100 epochs with the AdaBelief (Zhuang et al., 2020) optimizer with a batch size of 128 using PyTorch (Paszke et al., 2019).

²Note, that this is not explicitly stated in the paper, but can be observed in the original implementation <https://github.com/Philip-Bachman/Pseudo-Ensembles>.

The initial learning rate is set to 0.001 and halved every 25 epochs. We evaluate ensembles utilizing three different types of base models: Low capacity, mid-capacity, and large-capacity ones.

To do so, we use a ResNet architecture with 4 residual blocks, an input convolutional, and a linear layer for the output. All convolutions have a kernel size of 3×3 with padding and stride of one. They are always followed by a BatchNorm layer and ReLu activation. Each residual block consists of two convolutions and the residual connection followed by a 2×2 max pooling. In total each network has 9 convolution layers and a single linear layer. To vary the model capacity we use a different number of filters in each base model. The mid-capacity model utilizes 32 filters in each layer leading to 88 100 trainable parameters. The large-capacity model utilizes 96 filters leading to 706 468 trainable parameters. As the low-capacity model, we use a binarized version of the mid-capacity model which constrains the weights to $\{-1 + 1\}$. Binarized Neural Networks are a resource-friendly variation of ‘regular’ floating-point Neural Networks that are optimized towards minimal memory consumption and fast model application. They have been shown to perform nearly as good as their floating-point siblings while being magnitudes faster (Hubara et al., 2016; Rastegari et al., 2016; Zhu et al., 2019; Buschjäger et al., 2020). To train these models we use stochastic binarization, which retains the floating-point weights during the backward-pass but binarizes them during the forward pass as explained in (Hubara et al., 2016). We also evaluated EfficientNet-B0 (Tan & Le, 2019), MobilenetV3 (Howard et al., 2019) and DenseNets (Huang et al., 2017b) as base learners which yielded similar performance with more parameters and longer training times. We note, that EfficientNet-B0 and MobilenetV3 are optimized towards the larger images of ImageNet (Deng et al., 2009) (typically 224×224 pixels) and heavily downsample input images in the first layers. We hypothesize that this is not really necessary on the comparably smaller images of CIFAR100 leading to a similar performance with more parameters.

As discussed previously, GNCL encapsulates many existing methods and thus we compare GNCL to those methods not directly captured by it. We compare ensembles with $M = 16$ models trained via Bagging, via Stochastic Multiple Choice Learning (SMCL), via Gradient Boosting (GB), via Snapshot Ensembling (SE) and with Generalized Negative Correlation Learning (GNCL) minimizing the cross-entropy loss. For GNCL, we vary the regularization trade-off $\lambda \in \{0, 0.1, 0.2, \dots, 1.0\}$. Note, that GNCL with $\lambda = 0$ can be viewed as independent (Ind.) training of each network similar to Bagging but without bootstrap sampling. Similarly, for $\lambda = 1.0$ we train the ensemble in an End-To-End (E2E) fashion. For SE we take a snapshot of the model during optimization at the beginning of

epochs $\{2, 3, 4, 5, 10, 15, 20, 25, 30, 40, \dots, 90\}$ and combine them with the final model after 100 epochs. Last, we also train a single model for reference. Please note, that these experiments are not meant to produce benchmarking result on the CIFAR100 dataset, but to investigate the effects of ensembling with different algorithms and base learners.

Figure 1 depicts the results of this experiment. Each column depicts a different base learner with low, mid, and large capacity (from left to right). The first row depicts the test accuracy of the ensemble, the second row shows the average diversity among the ensemble members (evaluated on the test set), the third row shows the average accuracy of each expert model and the last row depicts the training loss. The horizontal bar depicts the test accuracy and training loss of a single model. Looking at the low capacity binarized neural networks, we see that they achieve an accuracy of around 5 – 60%. The clear winner in this setup GNCL with large $\lambda = 0.9$ and $\lambda = 1.0$ (E2E) achieving the highest accuracy, whereas SMCL is the worst with roughly 5% accuracy which is even below a single classifier. GNCL for smaller λ behaves similar to Bagging and Gradient Boosting and with larger λ there is a clear trend that the accuracy increases. Looking at the diversity we see that with larger λ it steeply increases while the average test accuracy expectantly decreases. Interestingly, SMCL offers a similar diversity to GNCL with $\lambda = 0.8$, but with much worse (average) test accuracy. The same effect can be observed for SE, but much less severe. The training loss indicates that a single model is under parameterized for the task at hand achieving a loss in the range of 2. Using more models increases the ensembles’ capacity, and therefore decreases the overall loss. Expectantly, the test accuracy roughly follows the training loss: The smaller the loss, the better the test accuracy where again E2E is the best. Looking at the mid-capacity models we see a similar picture as before, but note that the optimal test accuracy now shifts towards a smaller λ in the range of 0.8 – 0.9. Again, the diversity increases with increasing λ while the average test accuracy decreases. And again, we see that a smaller loss generally comes with better test accuracy. However, we note that while E2E learning offers by far the smallest loss it does *not* achieve the best test accuracy. This effect becomes more extreme when looking at large-capacity models. Here, GNCL with smaller λ in the range of 0 – 0.7 seems to be best, whereas for larger $\lambda > 0.7$ the performance reduces. As expected, the diversity increases with increasing λ while the average test accuracy decreases. Again we see, that E2E learning offers the smallest overall training loss, but it does *not* achieve the best test accuracy.

We conclude: For smaller capacity base models which do not achieve zero loss on their own, larger λ values and E2E learning seems to be best. In these cases, a large diversity *can* be beneficial as seen shown by the E2E approach. Once the base models become larger so that they achieve smaller




Figure 1. Experimental evaluation of various ensembling methods (E2E, Bagging, Gradient Boosting, GNCL with varying λ , SMCL and SE) with an ensemble size of $M = 16$ on CIFAR 100. Each column depicts a different base learner with low, mid, and large capacity (from left to right). The first row depicts the test accuracy of the ensemble, the second row shows the average diversity among the ensemble members (evaluated on test data), the third row shows the average accuracy of each expert and the last row depicts the training loss. The horizontal bar depicts the test accuracy and training loss of a single model.

losses on their own, enforcing diversity *can* be hurtful. This is clearly shown by SMCL which produces very diverse ensembles with sub-optimal performance but also shown by E2E learning for mid and large-capacity base learners. In this case, the training should shift towards a more independent training of each model with a smaller λ or independent training.

6. Conclusion

Ensemble learning plays a key role in many machine learning applications and offers state of the art performance. One of the guiding principles in designing an efficient ensembling algorithm is to enforce diversity in the ensemble. The theoretical roots of this approach lie in the bias-variance decomposition of the MSE loss which inspired many different approaches beyond the minimization of the MSE itself. While decompositions for other loss functions exist, they rarely inspired new learning algorithms beyond the general notion that diversity is important. In this paper, we studied the bias-variance decomposition for different loss functions more closely. We proposed a generalized bias-variance decomposition for twice differentiable loss functions which

implies that the diversity depends on the covariance of the experts’ outputs as well as the Hessian of the loss function. We derived a novel Generalized Negative Correlation Learning (GNCL) algorithm from it and detailed, how this algorithm encapsulates many existing works in literature. In an extensive experimental study we showed that diversity in the context of Deep Learning should not always be the main concern, but in fact, depends on the capacity of the base learners. For small-capacity base learners, diversity can be very beneficial as it allows the ensemble to minimize the overall loss more aggressively. For large-capacity base learners, diversity is also important but might be hurtful at a certain point as it artificially reduces the performance of the base learners and thus hurts their bias. This opens up the question, what Neural Network architectures are better suited for ensembling and how the optimization process impacts these results which we want to explore in the future.

References

Arora, S., Golowich, N., Cohen, N., and Hu, W. A convergence analysis of gradient descent for deep linear neural networks. *7th International Conference on Learn-*

- ing Representations, *ICLR 2019*, 2019. URL <https://arxiv.org/pdf/1810.02281.pdf>.
- Bachman, P., Alsharif, O., and Precup, D. Learning with pseudo-ensembles. In *Advances in Neural Information Processing Systems*, 2014.
- Baldi, P. and Sadowski, P. Understanding dropout. In *Advances in Neural Information Processing Systems*, 2013. doi: 10.17744/mehc.25.2.xhyreggxcd0q4ny.
- Bauer, E. and Kohavi, R. Empirical comparison of voting classification algorithms: bagging, boosting, and variants. *Machine Learning*, 1999. ISSN 08856125. doi: 10.1023/a:1007515423169.
- Biau, G. and Scornet, E. A Random Forest Guided Tour. *TEST*, 25(2):197–227, 2015. ISSN 1133-0686. doi: 10.1007/s11749-016-0481-7. URL <https://link.springer.com/content/pdf/10.1007/s11749-016-0481-7.pdf><http://arxiv.org/abs/1511.05741>.
- Breiman, L. Bagging predictors. *Machine Learning*, 1996. ISSN 08856125. doi: 10.1007/bf00058655.
- Brown, G., WatT, J. L., and Tino, P. Managing Diversity in Regression Ensembles. *Jmlr*, (6):1621–1650, 2005. ISSN 15505081. doi: 10.1097/IYC.000000000000008. URL <http://www.jmlr.org/papers/volume6/brown05a/brown05a.pdf>.
- Buschjäger, S., Pfahler, L., Buss, J., Morik, K., and Rhode, W. On-site gamma-hadron separation with deep learning on fpgas. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2020.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Devlin, J., Chang, M. W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1(Mlm): 4171–4186, 2019. URL <https://arxiv.org/pdf/1810.04805.pdf>.
- Domingos, P. A Unified Bias-Variance Decomposition. *Aaai/Iaai*, 2000.
- Dutt, A., Pellerin, D., and Quénot, G. Coupled ensembles of neural networks. *Neurocomputing*, 396:346–357, sep 2017. ISSN 18728286. doi: 10.1016/j.neucom.2018.10.092. URL <https://arxiv.org/abs/1709.06053>.
- Dvornik, N., Mairal, J., and Schmid, C. Diversity with co-operation: Ensemble methods for few-shot classification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. ISBN 9781728148038. doi: 10.1109/ICCV.2019.00382.
- Gal, Y. and Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *33rd International Conference on Machine Learning, ICML 2016*, 2016. ISBN 9781510829008.
- Geman, S., Bienenstock, E., and Doursat, R. Neural Networks and the Bias/Variance Dilemma, 1992. ISSN 0899-7667. URL <http://www.dam.brown.edu/people/geman/Homepage/Essaysandideasaboutneurobiology/bias-variance.pdf>.
- Geurts, P., Ernst, D., and Wehenkel, L. Extremely randomized trees. *Machine Learning*, 2006. ISSN 08856125. doi: 10.1007/s10994-006-6226-1.
- Hansen, J. V. and Heskes, T. General bias/variance decomposition with target independent variance of error functions derived from the exponential family of distributions. *Proceedings - International Conference on Pattern Recognition*, 15(2):207–210, 2000. ISSN 10514651. doi: 10.1109/icpr.2000.906049.
- Heskes, T. Selecting weighting factors in logarithmic opinion pools. In *Advances in Neural Information Processing Systems*, 1998. ISBN 0262100762.
- Howard, A., Sandler, M., Chen, B., Wang, W., Chen, L. C., Tan, M., Chu, G., Vasudevan, V., Zhu, Y., Pang, R., Le, Q., and Adam, H. Searching for mobileNetV3. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. ISBN 9781728148038. doi: 10.1109/ICCV.2019.00140.
- Huang, F., Ash, J. T., Langford, J., and Schapire, R. E. Learning deep ResNet blocks sequentially using boosting theory. In *35th International Conference on Machine Learning, ICML 2018*, 2018. ISBN 9781510867963.
- Huang, G., Li, Y., Pleiss, G., Liu, Z., Hopcroft, J. E., and Weinberger, K. Q. Snapshot ensembles: Train 1, get M for free. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, pp. 1–14, 2017a. URL <https://arxiv.org/pdf/1704.00109.pdf>.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings - 30th IEEE Conference on Computer Vision*

- and *Pattern Recognition, CVPR 2017*, 2017b. ISBN 9781538604571. doi: 10.1109/CVPR.2017.243.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Binarized neural networks. In *Advances in Neural Information Processing Systems*, 2016.
- James, G. M. Variance and bias for general loss functions. *Machine Learning*, 51(2):115–135, May 2003. ISSN 1573-0565. doi: 10.1023/A:1022899518027. URL <https://doi.org/10.1023/A:1022899518027>.
- Jiang, Z., Liu, H., Fu, B., and Wu, Z. Generalized ambiguity decompositions for classification with applications in active learning and unsupervised ensemble pruning. In *31st AAAI Conference on Artificial Intelligence, AAAI 2017*, 2017.
- Kawaguchi, K., Kaelbling, L. P., and Bengio, Y. Generalization in Deep Learning. 2017. URL <http://arxiv.org/abs/1710.05468>.
- Krizhevsky, A. Learning Multiple Layers of Features from Tiny Images. ... *Science Department, University of Toronto, Tech. ...*, 2009. ISSN 1098-6596. doi: 10.1.1.222.9220.
- Krogh, A. and Vedelsby, J. Neural Network Ensembles, Cross Validation, and Active Learning. *Advances in Neural Information Processing Systems 7*, 1995. ISSN 10495258.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 2017-December(Nips):6403–6414, 2017. ISSN 10495258. URL <https://arxiv.org/pdf/1612.01474.pdf>.
- Lee, S., Purushwalkam, S., Cogswell, M., Crandall, D., and Batra, D. Why M Heads are Better than One: Training a Diverse Ensemble of Deep Networks. 2015. URL <http://arxiv.org/abs/1511.06314>.
- Lee, S., Purushwalkam, S., Cogswell, M., Ranjan, V., Crandall, D., and Batra, D. Stochastic multiple choice learning for training diverse deep ensembles. *Advances in Neural Information Processing Systems*, 1(Nips):2127–2135, 2016. ISSN 10495258.
- Liu, Y. and Yao, X. Ensemble learning via negative correlation. *Neural Networks*, 12(10):1399–1404, 1999. ISSN 08936080. doi: 10.1016/S0893-6080(99)00073-8. URL <https://www.cs.bham.ac.uk/~pjt/NC/ncl.pdf>.
- Maclin, R. and Opitz, D. An empirical evaluation of bagging and boosting. *AAAI/IAAI*, 1997:546–551, 1997.
- Markowitz, H. The Utility of Wealth. *Journal of Political Economy*, 1952. ISSN 0022-3808. doi: 10.1086/257177.
- Mason, L., Baxter, J., Bartlett, P., and Frean, M. Boosting algorithms as gradient descent. In *Advances in Neural Information Processing Systems*, 2000. ISBN 0262194503.
- Melville, P. and Mooney, R. J. Creating diversity in ensembles using artificial data. *Information Fusion*, 2005. ISSN 15662535. doi: 10.1016/j.inffus.2004.04.001.
- Moghimi, M., Saberian, M., Yang, J., Li, L. J., Vasconcelos, N., and Belongie, S. Boosted convolutional neural networks. In *British Machine Vision Conference 2016, BMVC 2016*, 2016. doi: 10.5244/C.30.24.
- Opitz, D. W. and Maclin, R. F. An empirical evaluation of bagging and boosting for artificial neural networks. In *Proceedings of International Conference on Neural Networks (ICNN'97)*, volume 3, pp. 1401–1405 vol.3, 1997. doi: 10.1109/ICNN.1997.613999.
- Opitz, M., Possegger, H., and Bischof, H. Efficient model averaging for deep neural networks. *Asian Conference on Computer Vision*, 10112 LNCS:205–220, 2016. ISSN 16113349. doi: 10.1007/978-3-319-54184-6_13. URL https://www.tugraz.at/fileadmin/user/_upload/Institute/ICG/Documents/lrs/pubs/Opitz_accv_16.pdf.
- Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J. V., Lakshminarayanan, B., and Snoek, J. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing Systems*, 32(NeurIPS), 2019. ISSN 10495258. URL <https://arxiv.org/pdf/1906.02530.pdf>.
- Oza, N. C. and Russell, S. Online Bagging and Boosting. 2001. URL <https://ti.arc.nasa.gov/m/profile/oza/files/ozru01a.pdf>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.
- Qiu, X., Zhang, L., Ren, Y., Suganthan, P., and Amaratunga, G. Ensemble deep learning for regression and time series forecasting. *IEEE SSCI 2014 - 2014 IEEE Symposium Series on Computational Intelligence - CIEL 2014: 2014*

- IEEE Symposium on Computational Intelligence in Ensemble Learning, Proceedings*, (December), 2014. doi: 10.1109/CIEL.2014.7015739.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. XNOR-net: Imagenet classification using binary convolutional neural networks. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016. ISBN 9783319464923. doi: 10.1007/978-3-319-46493-0_32.
- Schapire, R. E. and Freund, Y. *Boosting: Foundations and algorithms*. MIT press, 2012.
- Schwenk, H. and Bengio, Y. Boosting neural networks. *Neural Computation*, 2000. ISSN 08997667. doi: 10.1162/089976600300015178.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 2014. ISSN 15337928.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. On the importance of initialization and momentum in deep learning. *30th International Conference on Machine Learning, ICML 2013, (PART 3)*:2176–2184, 2013. URL <http://proceedings.mlr.press/v28/sutskever13.pdf>.
- Tan, M. and Le, Q. V. EfficientNet: Rethinking model scaling for convolutional neural networks. In *36th International Conference on Machine Learning, ICML 2019*, 2019. ISBN 9781510886988.
- Ueda, N. and Nakano, R. Generalization error of ensemble estimators. *Proceedings of International Conference on Neural Networks (ICNN'96)*, 1(xi):90–95, 1996. doi: 10.1109/ICNN.1996.548872. URL <http://ieeexplore.ieee.org/document/548872/>.
- Webb, A. M., Reynolds, C., Iliescu, D.-A., Reeve, H., Lujan, M., and Brown, G. Joint Training of Neural Network Ensembles. (4):1–14, 2019. doi: 10.13140/RG.2.2.28091.46880. URL <http://arxiv.org/abs/1902.04422>{%}0A<http://dx.doi.org/10.13140/RG.2.2.28091.46880>.
- Webb, A. M., Reynolds, C., Chen, W., Reeve, H., Iliescu, D.-A., Lujan, M., and Brown, G. To Ensemble or Not Ensemble: When does End-To-End Training Fail? (2018):1–16, 2020. doi: 10.13140/RG.2.2.28091.46880. URL <http://arxiv.org/abs/1902.04422>{%}0A<http://dx.doi.org/10.13140/RG.2.2.28091.46880>.
- Webb, G. I. MultiBoosting: a technique for combining boosting and wagging. *Machine Learning*, 2000. ISSN 08856125. doi: 10.1023/A:1007659514849.
- Zeiler, M. D. and Fergus, R. Visualizing and understanding convolutional networks. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2014. ISBN 9783319105895. doi: 10.1007/978-3-319-10590-1_53.
- Zhang, C., Recht, B., Bengio, S., Hardt, M., and Vinyals, O. Understanding deep learning requires rethinking generalization. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, 2017. URL <https://arxiv.org/pdf/1611.03530.pdf?from=timeline{%}isappinstalled=0>.
- Zhou, Z.-H. and Feng, J. Deep Forest. feb 2017. URL <http://arxiv.org/abs/1702.08835>.
- Zhu, S., Dong, X., and Su, H. Binary ensemble neural network: More bits per network or more networks per bit? In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019. ISBN 9781728132938. doi: 10.1109/CVPR.2019.00506.
- Zhuang, J., Tang, T., Ding, Y., Tatikonda, S., Dvornek, N., Papademetris, X., and Duncan, J. Adabelief optimizer: Adapting stepsizes by the belief in observed gradients. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.