

A Unified Theory of Diversity in Ensemble Learning

Danny Wood

Tingting Mu

Andrew M. Webb

Henry W. J. Reeve

Mikel Luján

Gavin Brown

DANNY.WOOD@MANCHESTER.AC.UK[†]

TINGTING.MU@MANCHESTER.AC.UK[†]

ANDREW.WEBB@MANCHESTER.AC.UK[†]

HENRY.REEVE@BRISTOL.AC.UK*

MIKEL.LUJAN@MANCHESTER.AC.UK[†]

GAVIN.BROWN@MANCHESTER.AC.UK[†]

[†] *Department of Computer Science, University of Manchester, UK*

* *School of Mathematics, University of Bristol, UK*

Editor: Boaz Nadler

Abstract

We present a theory of ensemble diversity, explaining the nature of diversity for a wide range of supervised learning scenarios. This challenge has been referred to as the “holy grail” of ensemble learning, an open research issue for over 30 years. Our framework reveals that diversity is in fact a *hidden dimension* in the bias-variance decomposition of the ensemble loss. We prove a family of *exact* bias-variance-diversity decompositions, for a wide range of losses in both regression and classification, e.g., squared, cross-entropy, and Poisson losses. For losses where an additive bias-variance decomposition is not available (e.g., 0/1 loss) we present an alternative approach: quantifying the *effects* of diversity, which turn out to be dependent on the label distribution. Overall, *we argue that diversity is a measure of model fit, in precisely the same sense as bias and variance*, but accounting for statistical dependencies *between* ensemble members. Thus, we should not be ‘maximising diversity’ as so many works aim to do—instead, we have a bias/variance/diversity trade-off to manage.

Keywords: ensembles, diversity, bias, variance

1. Introduction

Ensemble methods have enabled state-of-the-art results for decades: from early industrial computer vision (Viola and Jones, 2001) to the deep learning revolution (Krizhevsky et al., 2012), and inter-disciplinary applications (Cao et al., 2020). An accepted mantra is that ensembles work best when the individuals have a “diversity” of predictions—often induced by classical methods such as Bagging (Breiman, 1996), but diversity-encouraging heuristics are rife in the literature (Brown et al., 2005). Given this, we trust that the ensemble will “average out” the errors of the individuals. One reason for the popularity of such methods is clear: the very *idea* of ensembles is an appealing anthropomorphism, invoking analogies to human committees, and “wisdom of the crowds”. However, such analogies have limitations. More formal approaches have been pursued, in particular for *quantifying* diversity. It is obvious that we do not want all predictions to be identical; and, it is equally obvious we do not want them to be different just for the sake of it, sacrificing overall performance. We want something in-between these two—the so-called

accuracy/diversity trade-off. However, here we encounter the problem of formally defining “diversity” and its relation to ensemble performance. *In general, there is no agreement on how to quantify diversity*, except in the limited case of regression with an arithmetic mean ensemble (Krogh and Vedelsby, 1994; Ueda and Nakano, 1996). For classification and other scenarios, there are dozens of proposed diversity measures (Kuncheva, 2014). A comprehensive theory of ensemble diversity has been an open problem for over 30 years.

Motivation: Our primary motivation is to fill this ‘gap’ in current ensemble theory, providing a solid foundation to understand and study ensemble diversity. However, there are also many practical reasons to pursue this. Diverse ensembles can be more computationally efficient than single large models, with the *same* generalisation performance (Kondratyuk et al., 2020). Diverse ensembles are robust against adversarial attacks (Biggio et al., 2011; Pang et al., 2019), and can counteract covariate shift (Sinha et al., 2020). Advantages are also found in important application areas (Cao et al., 2020) and well beyond supervised learning (Carreira-Perpinán and Raziperchikolaei, 2016). It is important to note that these use-cases *do not follow a common approach*: they either adopt some measurement of diversity picked from historical literature, or propose their own novel metric. There is, therefore, reason to pursue a unified theory, where diversity is derived from first principles.

This challenge has proven non-trivial: surveys of progress can be found in Dietterich (2000); Brown et al. (2005); Zhou (2012), and Kuncheva (2014). Diversity is nowadays referred to as a heuristic with no precise definition, and, it has been said:

“There is no doubt that understanding diversity is the holy grail in the field of ensemble learning” (Zhou, 2012, Sec 5.1, page 100).

Summary of our Results: In contrast to previous efforts which *define* novel diversity measures, we take loss functions and *decompose* them, exposing terms that naturally account for diversity. We show that diversity is a *hidden dimension* in the *bias-variance decomposition* of the ensemble loss. In particular, we prove *exact* bias-variance-*diversity* decompositions, applying for a broad range of losses, taking a common form:

$$\text{expected loss} = (\text{average bias}) + (\text{average variance}) - (\text{diversity}),$$

where diversity is a measure of ensemble member disagreement, independent of the label. For the special case of squared loss, this is an alternative to a decomposition proposed by Ueda and Nakano (1996), but expanding the formal notion of diversity to many other losses, e.g., the *cross-entropy*, and the Poisson loss. For losses where an additive bias-variance decomposition does not exist (e.g., 0/1 loss) we present an approach to precisely quantify the *effects* of diversity, with the caveat that the effects are conditional on the label distribution.

Overall, *we argue that diversity is best understood as a measure of model fit, in precisely the same sense as bias and variance*, but accounting for statistical dependencies *between* ensemble members. Thus, we should not be ‘maximising’ diversity as so many works claim to do—instead, we have a trade-off to manage. With single models, we have the well-known *bias/variance* trade-off. With an ensemble we have a *bias/variance/diversity* trade-off, varying both with individual model capacity, *and* similarities between model predictions.

Notation

For quick reference, we summarise the majority of notational conventions used in the paper. More detailed definitions are given in the main body as appropriate. Boldface font, e.g., \mathbf{y} , denotes a vector, whereas y is a scalar. Capital letters, e.g., Y , denote a random variable, where y is a realisation of Y .

For brevity, in the majority of the paper we leave it implicit that a model, q , is a function of an input \mathbf{x} , and is dependent on a training set sampled from a random variable D .

<u>Symbol</u>	<u>Description</u>
$\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$	Input, in d dimensions
$\mathbf{y} \in \mathcal{Y} \subseteq \mathbb{R}^k$	Label, in k dimensions (or y for a scalar)
$\{(\mathbf{x}_j, \mathbf{y}_j)\}_{j=1}^n$	Training set of n labelled examples
$P(\mathbf{x}, y)$	Unknown data distribution over (\mathbf{x}, y)
$\mathbb{E}_{\mathbf{X}Y}[\dots]$	Expectation w/r $P(\mathbf{x}, y)$, i.e., over all possible test examples (\mathbf{x}, y)
	$\mathbb{E}_{\mathbf{X}Y}[\dots] \stackrel{\text{def}}{=} \int P_{\mathbf{X}}(\mathbf{x}) \int P_{Y \mathbf{X}}(y \mathbf{x})[\dots] dy d\mathbf{x}$
D	Random variable $D \sim P(\mathbf{x}, y)^n$, over i.i.d. training sets of size n
$\mathbb{E}_D[\dots]$	Expectation over all possible training sets $\{(\mathbf{x}_j, \mathbf{y}_j)\}_{j=1}^n$ drawn from D
$q(\mathbf{x})$	Predictive model q , mapping $\mathbf{x} \rightarrow y$. If output is a vector then $\mathbf{q}(\mathbf{x})$
$\ell(y, q(\mathbf{x}))$	Loss function for label y and prediction $q(\mathbf{x})$
$\ell_{0/1}(y, q(\mathbf{x}))$	The 0/1 loss function: returns 1 if $y \neq q(\mathbf{x})$, and 0 otherwise.
$R(q)$	Risk of model q , defined $R(q) := \mathbb{E}_{\mathbf{X}Y}[\ell(y, q(\mathbf{x}))]$
m	Number of ensemble members
q_i	The i th member of an ensemble, $\{q_i\}_{i=1}^m$
\bar{q}	Centroid combiner rule, defined $\bar{q} := \arg \min_{z \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m [\ell(z, q_i)]$
\hat{q}	Centroid of a model distribution w/r D , defined $\hat{q} := \arg \min_{z \in \mathcal{Y}} \mathbb{E}_D[\ell(z, q)]$
$\phi(\mathbf{y})$	Strictly convex generator function for a Bregman divergence
$\nabla \phi(\mathbf{y})$	Gradient vector of generator ϕ w/r argument \mathbf{y}
$B_\phi(\mathbf{y}, \mathbf{q})$	Bregman divergence between points $\mathbf{y}, \mathbf{q} \in \mathbb{R}^k$, using generator ϕ

Table 1: Notation used in the paper.

2. Problem Statement: What is Ensemble Diversity?

Hansen and Salamon (1990) proposed a methodology to train multiple neural networks, encouraging different models by providing each with a different training data subset. Many subsequent papers followed this “parallel” strategy: an early work being Perrone and Cooper (1992), but perhaps more well-known are *Bagging* (Breiman, 1996), and *Random Forests* (Breiman, 2001). *Boosting* algorithms (Schapire et al., 1998) exploit a similar principle, but construct models *sequentially*. These approaches, parallel and sequential (see Figure 1), are the most common schemes to learn ensembles (Kuncheva, 2014), and there are many variations, e.g., selecting members from a pre-constructed candidate pool, each built by independent (parallel) teams working on the same task but with different data.

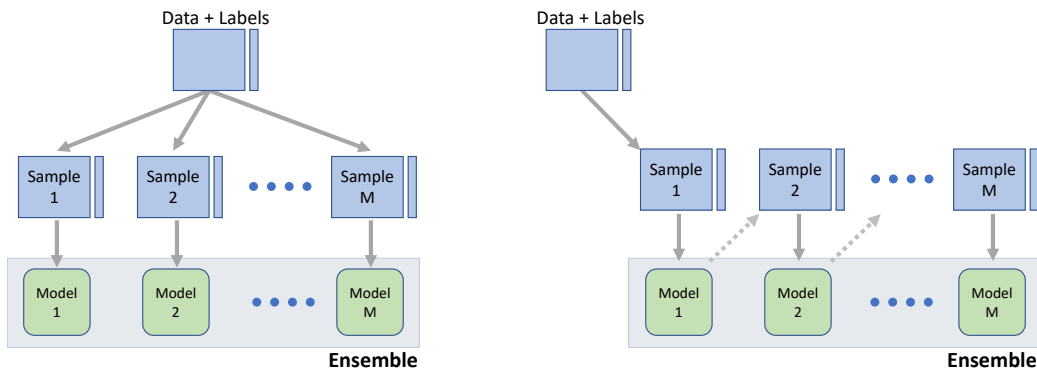


Figure 1: Parallel vs sequential ensemble construction. Both can be seen as creating “diverse” models in some sense—either *implicitly* (independently re-sampling the training data), or *explicitly* (re-sampling according to the errors of earlier models).

So why do these strategies work? Both can be understood heuristically in terms of “diversity”, in the sense coined by Opitz and Shavlik (1996), referring to differences in generalisation behavior among a group of models. In a review, Dietterich (2000) explains:

“An accurate classifier is one that has an error rate of better than random guessing on new x values. Two classifiers are diverse if they make different errors on new data points.” (Dietterich, 2000)

In this sense, both approaches foster diversity—either *implicitly* by randomly perturbing the data for each model, or *explicitly* by constructing each data set to address the errors of previous models (Brown et al., 2005). The implicit approach has been widely adopted in deep learning: Goodfellow et al. (2016) note that sources of randomness in the initialisation of deep networks are often enough to cause them to make partially independent errors. Given the success of ensembles, there have been many attempts to explain *why* they work, in terms of the diversity. Goodfellow et al. (2016) writes,

“The reason that model averaging works is that different models will usually not make all the same errors on the test set”. (Goodfellow et al., 2016, p. 249)

Whilst this is true, we desire a more formal treatment.

What are we looking for? A theory of diversity would ideally have three key ingredients:

1. a definition of diversity as a measure of disagreement between the ensemble members;
2. this measure should have a clear relation to the overall ensemble error; and,
3. the theory should have a clear relation to previously established results, and expand our understanding to a wider range of learning scenarios.

The first point concerns the definition of diversity itself. An interesting question is whether the measure is a function solely of the ensemble member outputs, or whether it can involve the label distribution as well. Kuncheva (2014) refers to measures that rely on the true label as ‘oracle’ diversity measures. *Both outcomes would be an interesting scientific conclusion:* the label-independent case showing that diversity is solely a function of the ensemble itself, and the label-dependent case showing that it is a class-conditional phenomenon.

The second point ensures we can interpret what effect diversity has on our ultimate objective: reducing the ensemble error. The third point ensures that the new theory contributes to knowledge, in the sense that it expands the depth/breadth of our understanding. This also relates to the only known scenario where the challenge of defining diversity might be considered a “solved” problem: regression using squared loss, with an arithmetic mean ensemble. We now review this.

Known results for regression ensembles: Krogh and Vedelsby (1994) showed that, for an arithmetic mean combiner, using squared loss, the ensemble loss *is guaranteed to be less than or equal to the average individual loss.*

Theorem 1 (Ambiguity decomposition, Krogh & Vedelsby, 1994) *Given a label $y \in \mathbb{R}$, a member prediction $q_i(\mathbf{x})$, and an ensemble $\bar{q}(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m q_i(\mathbf{x})$, we have,*

$$(\bar{q}(\mathbf{x}) - y)^2 = \frac{1}{m} \sum_{i=1}^m (q_i(\mathbf{x}) - y)^2 - \frac{1}{m} \sum_{i=1}^m (q_i(\mathbf{x}) - \bar{q}(\mathbf{x}))^2. \quad (1)$$

The left hand side is the ensemble loss at a single test point (\mathbf{x}, y) . The first term on the right is the average individual loss. The second is known as the *ambiguity*—measuring the disagreement of individuals, as a spread around the ensemble prediction. Since this is non-negative, it guarantees the ensemble loss will be less than or equal to the average loss.

This result is often *erroneously* cited as the reason why all ensembles work. However, the form above applies *only if we use the squared loss with an arithmetic mean combiner.* If we use squared loss with a different combiner, the result no longer holds. A deeper understanding came from Ueda and Nakano (1996)—though under the same loss/combiner assumptions. They extended the classical bias-variance theory of Geman et al. (1992) to show that the *expected* squared loss of the ensemble decomposes into three terms, involving the bias, variance, and *covariance* of the models. Here, the covariance captures the notion of ‘diversity’. More details on this will be provided later.

As mentioned, the expression above does not apply beyond squared loss with the arithmetic mean combiner. A significant community effort has been directed to find corresponding notions of diversity for classification problems. We review this next.

Known results for classifier ensembles: For classification problems, we might have estimates of the class probability distribution, or just labels—usually combined by averaging probabilities, or majority voting. Tumer and Ghosh (1996) showed that the correlation between pairs of averaged class probabilities had a simple relationship to the overall ensemble classification error, at least in a region close to decision boundaries. Brown (2009) and Zhou and Li (2010) proposed information theoretic analyses, showing that diversity manifests as both low- and high-order interactions between ensemble members. Buschjäger et al. (2020) used a Taylor approximation on twice-differentiable losses, showing an exact decomposition only when higher derivatives are zero, e.g., squared loss, but not cross-entropy. Similarly, Ortega et al. (2022) decomposed upper bounds on losses, again only obtaining an equality for squared loss. Kuncheva and Whitaker (2003) considered diversity measures for their correlation to the ensemble error. They investigated numerous discrepancy metrics in the form $\delta(q_i(\mathbf{x}), q_j(\mathbf{x})) \in \mathbb{R}$, which are averaged over all pairs of ensemble members at each test point \mathbf{x} :

$$\text{diversity}(q_1, \dots, q_m, \mathbf{x}) = \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j \neq i} \delta(q_i(\mathbf{x}), q_j(\mathbf{x})). \quad (2)$$

The ensemble diversity is this quantity averaged over a validation dataset. The ensemble diversity is evaluated on its empirical correlation to the ensemble performance, and seen as more successful if it has high correlation, illustrated in Figure 2.




Figure 2: Accuracy/diversity for two (hypothetical) diversity measures. Measure B (right) is more desirable, as it has stronger correlation to performance improvement.

Several measures (including also non-pairwise measures) were explored, with no single measure proving more successful than any other. Almost 20 years on, novel diversity heuristics/measures are still being proposed, e.g., Jan and Verma (2019); Wu et al. (2021).

Our approach to the problem: As above, diversity is often discussed in a dichotomy of classification vs. regression. One of our primary observations is that *this high-level dichotomy is insufficient to fully describe the nature of diversity*. Instead, we must consider the *loss function at hand*. For regression, this might be squared loss, but also Poisson, or others. For classification, this may be 0/1, or cross-entropy. We build on the strong foundation of bias-variance theory (Geman et al., 1992; James and Hastie, 1997; Heskes, 1998; Pfau, 2013), and *decompose* the losses, naturally exposing notions of diversity.

3. A Very Short Introduction to Bias-Variance Decompositions

We review literature on *bias-variance* decompositions: the statistical setting is a standard supervised learning scenario. We assume a training set $\{(\mathbf{x}_j, y_j)\}_{j=1}^n$ drawn from a random variable $D \sim P(\mathbf{X}, Y)^n$. From this, we learn a model, q . For brevity (see Table 1) in the remainder of the paper we leave it implicit that q is a function of \mathbf{x} , and is always dependent on data sampled from D . We denote the *risk* as $R(q) \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{X}Y} [\ell(Y, q)]$, where the label Y is a random variable, thus potentially noisy. If $\ell(y, q) = (y - q)^2$, this is the *squared risk*, and the Bayes-optimal prediction is the conditional mean $Y^* \stackrel{\text{def}}{=} \mathbb{E}_{Y|\mathbf{X}}[Y]$. Geman et al. (1992) considered the *expected squared risk*, over possible training sets drawn from D , showing the following three term decomposition.

$$\underbrace{\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}Y} [(Y - q)^2] \right]}_{\text{expected squared risk}} = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{Y|\mathbf{X}} [(Y - Y^*)^2]}_{\text{noise}} + \underbrace{\left(\mathbb{E}_D [q] - Y^* \right)^2}_{\text{bias}} + \underbrace{\mathbb{E}_D \left[\left(q - \mathbb{E}_D [q] \right)^2 \right]}_{\text{variance}} \right]. \quad (3)$$

The first term is the irreducible *noise* in the problem, independent of any model parameters. The second is the *bias*¹, defined as the loss of the expected model against the Bayes-optimal prediction. The third is the *variance*, capturing variation in q due to different training sets. Note that we can use the term *bias* (correspondingly *variance*) to indicate the value at a point (\mathbf{x}, y) , or averaged over a distribution—the intention will be clear from context. The ideas are often explained with a dartboard diagram, as in Figure 3.




Figure 3: The classic dartboard analogy for explaining bias and variance.

The bullseye (yellow circle) is the target for a single *test* point, and each blue dot is a prediction from a model given a different *training* set. A model with high bias/low variance will be on average far from the target, and insensitive to small data changes. A model with low bias/high variance will have an expected value close to the target, but be very sensitive to data changes, meaning any given model is likely to overfit. The random variable D does not *have* to be over data, but can be any stochastic quantity involved in learning the model, e.g., random weights in a neural network.

1. *Is it ‘squared bias’, or just ‘bias’?* Geman et al’s bias term is often referred to as “squared bias”. However, the square is in fact an artefact of using the squared loss, and is not present in the decompositions of other losses. Thus, the term “squared bias” is a *misnomer*, and throughout the paper we use simply “bias”. Further exposition on this point is provided in Appendix B.1.

Bias-variance decompositions apply for more than just squared loss. Geman et al. (1992) is a widely appreciated result. It is less well known that similar decompositions hold for other losses, e.g., the KL-divergence of class probability estimates (Heskes, 1998). Assume \mathbf{y} is a one-hot vector of length k , and $\mathbf{q} \in \mathbb{R}^k$ is our model’s predicted distribution over the k labels. The expected risk decomposes² in a similar three-part form:

$$\underbrace{\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} [K(\mathbf{Y} \parallel \mathbf{q})] \right]}_{\text{expected risk}} = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [K(\mathbf{Y} \parallel \mathbf{Y}^*)]}_{\text{noise}} + \underbrace{K(\mathbf{Y}^* \parallel \hat{\mathbf{q}})}_{\text{bias}} + \underbrace{\mathbb{E}_D [K(\hat{\mathbf{q}} \parallel \mathbf{q})]}_{\text{variance}} \right], \quad (4)$$

where $\hat{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D [K(\mathbf{z} \parallel \mathbf{q})] = Z^{-1} \exp(\mathbb{E}_D[\ln \mathbf{q}])$ is a *normalized geometric mean* at each input vector \mathbf{x} . This can also be illustrated as a “dartboard” analogy, in Figure 4.




Figure 4: Dartboard diagram illustrating bias/variance for the KL-divergence.

Bias-variance decompositions do not apply for all losses. A commonality of the decompositions above is that the variance is *independent* of the label. A further commonality is that the bias is the loss of a predictor $\hat{q} \stackrel{\text{def}}{=} \arg \min_{z \in \mathcal{Y}} \mathbb{E}_D [\ell(z, q)]$, not dependent on any particular training set. We refer to this form as an *additive* bias-variance decomposition. It turns out that this does not hold for all losses, e.g., 0/1 loss—in this case \hat{q} is the *mode* of the random variable, but, we have an *inequality*:

$$\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} [\ell_{0/1}(Y, q)] \right] \neq \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [\ell_{0/1}(Y, Y^*)] + \ell_{0/1}(Y^*, \hat{q}) + \mathbb{E}_D [\ell_{0/1}(\hat{q}, q)] \right]. \quad (5)$$

Many authors have presented alternative 0/1 decompositions (Kohavi et al., 1996; Tibshirani, 1996; James and Hastie, 1997; Heskes, 1998; Domingos, 2000). An excellent review can be found in Geurts (2002). A common theme is to either sacrifice the simple additive form, or have a formulation of variance that is dependent on the label distribution. The necessary and sufficient conditions for an additive bias-variance decomposition of the form mentioned above, are an open research question.

Together, the theories outlined above have been used to understand the nature of model fitting—the so-called ‘bias-variance’ trade-off. In the following section we build upon all these theories to understand the phenomenon of ensemble diversity.

2. Heskes (1998) presents a slight variation, though the expressions are equivalent. See Appendix B.3.

4. A Unified Theory of Ensemble Diversity

Ensemble “diversity” is a popular, but variously defined idea. Bias and variance are clear-cut, and precisely defined. It makes sense to build stronger bridges between the two. Our approach is exactly this, revealing diversity as a *hidden dimension* in the bias-variance decomposition of an ensemble. We argue that diversity can be considered in exactly the same way as bias/variance—simply another quantifiable aspect of model fitting.

4.1 Diversity as a Hidden Dimension of the Bias-Variance Decomposition

Our approach yields a ‘unified theory’ in the sense that the *same* methodology applies for many different losses, in classification/regression scenarios. We first define a *generalised* (additive) bias-variance decomposition, in that it generalises Equations (3) and (4).

Definition 2 (Generalised Bias-Variance Decomposition) *For a loss ℓ , we define $\hat{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D[\ell(\mathbf{z}, \mathbf{q})]$ as the ‘centroid’ of the model distribution. If the following form holds, we refer to it as a generalised bias-variance decomposition.*

$$\underbrace{\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}}[\ell(\mathbf{Y}, \mathbf{q})] \right]}_{\text{expected risk}} = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\ell(\mathbf{Y}, \mathbf{Y}^*)]}_{\text{noise}} + \underbrace{\ell(\mathbf{Y}^*, \hat{\mathbf{q}})}_{\text{bias}} + \underbrace{\mathbb{E}_D[\ell(\hat{\mathbf{q}}, \mathbf{q})]}_{\text{variance}} \right]. \quad (6)$$

When $\ell(y, q) = (y - q)^2$, we have $\hat{q} = \mathbb{E}_D[q]$ and this reduces to Equation (3). Alternatively, with $\ell(\mathbf{y}, \mathbf{q}) = K(\mathbf{y} \parallel \mathbf{q})$, we have $\hat{\mathbf{q}} = Z^{-1} \exp(\mathbb{E}_D[\ln \mathbf{q}])$, and it reduces to Equation (4). Importantly, this same additive form also holds for the even broader family of *Bregman divergences* (Bregman, 1967). This family, reviewed in the next section, provides us a convenient analytical form, which we will use to derive several interesting properties.

We now take our first step toward ensembles. We highlight that the *ambiguity* decomposition, Equation (1), can be proven by a very similar method to the bias-variance decomposition. In fact: if one exists, then the other must also, outlined in Appendix B.2. For a loss ℓ , if Equation (6) holds, we can state the corresponding ambiguity decomposition.

Proposition 3 (Generalised Ambiguity Decomposition) *Assuming loss ℓ admits a bias-variance decomposition then, for an ensemble $\{\mathbf{q}_i\}_{i=1}^m$, the ambiguity decomposition is,*

$$\underbrace{\ell(\mathbf{y}, \bar{\mathbf{q}})}_{\text{ensemble loss}} = \underbrace{\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{y}, \mathbf{q}_i)}_{\text{average loss}} - \underbrace{\frac{1}{m} \sum_{i=1}^m \ell(\bar{\mathbf{q}}, \mathbf{q}_i)}_{\text{ambiguity}}, \quad (7)$$

where $\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i)$ is the ensemble combination.

We highlight that the ensemble combiner is defined **as a property of the loss**. For squared loss, this results in $\bar{q} = \frac{1}{m} \sum_{i=1}^m q_i$, the commonly used arithmetic mean combiner. For KL-divergence, this results in $\bar{\mathbf{q}} = Z^{-1} \prod_{i=1}^m \mathbf{q}_i^{1/m}$, i.e., a normalised geometric mean, also known as the ‘product rule’ (Kuncheva, 2014). Audhkhasi et al. (2013) and Jiang et al. (2017) proposed ambiguity decompositions for classification losses, though both assumed the combiner must be an arithmetic mean. In contrast, here the combiner is defined *in terms of the loss ℓ* . We refer to this as the *centroid combiner rule*, defined formally below.

Definition 4 (Centroid Combiner rule) For a point (\mathbf{x}, \mathbf{y}) , given a set of predictions $\{\mathbf{q}_i\}_{i=1}^m$, the centroid combiner $\bar{\mathbf{q}}$ is the minimizer of the averaged loss $\ell(\mathbf{z}, \mathbf{q}_i)$, over all ensemble members.

$$\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i). \quad (8)$$

For each given loss function, the centroid combiner *could* be derived separately. However, in the next section we will meet the family of *Bregman divergences* (Bregman, 1967), which provide us a common analytic form for the centroid combiner. We can now present our main result, Theorem 5, a generalised bias-variance-*diversity* decomposition.

Theorem 5 (Generalized Bias-Variance-Diversity decomposition) Consider a set of models $\{\mathbf{q}_i\}_{i=1}^m$, evaluated by a loss ℓ . Assuming a bias-variance decomposition holds in the form of Definition 2, the following decomposition also holds.

$$\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} [\ell(\mathbf{Y}, \bar{\mathbf{q}})] \right] = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [\ell(\mathbf{Y}, \mathbf{Y}^*)]}_{\text{noise}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{Y}^*, \hat{\mathbf{q}}_i)}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [\ell(\hat{\mathbf{q}}_i, \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m \ell(\bar{\mathbf{q}}, \mathbf{q}_i) \right]}_{\text{diversity}} \right], \quad (9)$$

where $\hat{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D [\ell(\mathbf{z}, \mathbf{q})]$ and the combiner is $\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i)$.

This has decomposed the expected ensemble risk into: the noise, the average bias, the average variance, and the expected ambiguity. It is this expected ambiguity term that we consider as the ensemble *diversity*, which we highlight has the opposite sign to other terms. We can make two important observations.

The ensemble combiner $\bar{\mathbf{q}}$ is a centroid. In a practical scenario, the loss of the ensemble is chosen for whatever the task is: 0/1 loss, cross-entropy, or MSE, or many others. The combiner rule can also be chosen as almost anything: voting, arithmetic mean, weighted combinations, geometric mean, etc. The result above shows that, if we impose a constraint between these choices (i.e., Definition 4) we enable a decomposition which naturally exposes a diversity term. *We emphasize that using this makes no claim on the empirical behavior of the ensemble, but simply on their theoretical properties.*

There is now a bias/variance/diversity trade-off. As individual models increase in capacity, their average bias decreases. Without regularisation, their average variance would increase. However, these terms determine only part of the ensemble behavior. The final part is the diversity. A critical point here is *diversity always subtracts from the expected risk*. This is not to say that greater diversity always reduces expected risk—it only reduces it given a fixed bias and variance. Ultimately, the three-way trade-off of *bias/variance/diversity* is what determines the overall ensemble performance. It is worth highlighting that diversity is defined similarly to bias/variance, involving an expectation over D , as opposed to being a property of a single training run. We also highlight the result applies for both dependent and independent training (where D is a vector of m random variables, see Appendix C.1 for details) and also for schemes where D defines a pool of pre-constructed models.

4.2 What if a bias-variance decomposition doesn't hold?

Definition 2 does not apply for all losses, e.g., for 0/1 loss, or absolute loss. This creates an obvious challenge for our framework. We relied on the existence of such a decomposition, deriving diversity as a ‘hidden dimension’. However, there is still a way forward.

Quantifying the Effect of Bias/Variance/Diversity. James and Hastie (1997) present a decomposition which applies for *any* loss, and retains the simple additive form. They achieve this by distinguishing the *measurement* of variance (which is independent of the label) from its *effect* on the expected risk (which is dependent on the label), with the analogous pair of quantities for the bias. The measurement and the effect are not necessarily the same quantity. They considered the *measurement* to be the ‘natural’ form in Definition 2. They then defined the *bias-effect*³, and the *variance-effect*. For compact notation, we average these over $P(\mathbf{x}, \mathbf{y})$. Using $R(\mathbf{q}) \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{X}\mathbf{Y}}[\ell(\mathbf{y}, \mathbf{q})]$ as the risk of \mathbf{q} , their terms are:

$$\text{bias-effect} \stackrel{\text{def}}{=} R(\mathring{\mathbf{q}}) - R(\mathbf{y}^*), \tag{10}$$

$$\text{variance-effect} \stackrel{\text{def}}{=} \mathbb{E}_D \left[R(\mathbf{q}) - R(\mathring{\mathbf{q}}) \right], \tag{11}$$

where $\mathring{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D \left[\ell(\mathbf{z}, \mathbf{q}) \right]$, and $\mathbf{y}^* \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} \left[\ell(\mathbf{y}, \mathbf{z}) \right]$ is the Bayes-optimal predictor⁴ at each point \mathbf{x} . These quantify the *effect on the risk* for using one predictor instead of another. The *bias-effect* is the change in risk, for the centroid $\mathring{\mathbf{q}}$ versus the Bayes-optimal predictor \mathbf{y}^* . The *variance-effect* is the change in risk for a model \mathbf{q} versus the centroid, averaged over the distribution of D . For losses where Definition 2 holds, the measurement and the effect are the same, i.e., variance-effect is equal to variance, and bias-effect is equal to bias. For example, with squared loss and $y \in \mathbb{R}$, Equation (10) is equal to $\mathbb{E}_{\mathbf{X}\mathbf{Y}}[(\mathbb{E}_D[q] - y^*)^2]$. For general losses (e.g., 0/1 loss, absolute loss), this is not the case, and the *effect* terms are different numerical quantities. With these definitions, James & Hastie note the following decomposition.

Theorem 6 (Bias-Variance Effect decomposition, James & Hastie 1997) *For a loss $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, denoting the centroid as $\mathring{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D [\ell(\mathbf{z}, \mathbf{q})]$, the expected risk decomposes as follows:*

$$\mathbb{E}_D \left[R(\mathbf{q}) \right] = \underbrace{R(\mathbf{y}^*)}_{\text{noise}} + \underbrace{R(\mathring{\mathbf{q}}) - R(\mathbf{y}^*)}_{\text{bias-effect}} + \underbrace{\mathbb{E}_D [R(\mathbf{q}) - R(\mathring{\mathbf{q}})]}_{\text{variance-effect}} \tag{12}$$

This decomposition holds for any loss, since terms on the right always cancel to give the left side. A notable point here is that the variance-effect is dependent on the label y , whilst the variance itself is a label-independent quantity. Furthermore, whilst bias-effect is non-negative, the variance-effect can be a *negative* quantity. From this, using the same relation exploited for Proposition 3, we can state a corresponding ambiguity (effect) decomposition.

3. J&H called this the *systematic* effect. For consistency within our work, we prefer the term *bias-effect*.
 4. We acknowledge a slight abuse of notation, denoting the Bayes predictor \mathbf{y}^* either as a function at each point \mathbf{x} , or a vector in \mathbb{R}^k , as needed. The intention will be made clear from context.

Proposition 7 (Ambiguity-Effect Decomposition) *Given a loss $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, and an ensemble $\{\mathbf{q}_i\}_{i=1}^m$ using the centroid combiner $\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \left[\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i) \right]$, we have a decomposition:*

$$\underbrace{R(\bar{\mathbf{q}})}_{\text{ensemble risk}} = \underbrace{\frac{1}{m} \sum_{i=1}^m R(\mathbf{q}_i)}_{\text{average risk}} - \underbrace{\frac{1}{m} \sum_{i=1}^m \left[R(\mathbf{q}_i) - R(\bar{\mathbf{q}}) \right]}_{\text{ambiguity-effect}}. \quad (13)$$

We re-emphasize that this applies for any loss. This is again quite obvious, as terms on the right cancel to give the left hand side. For losses where Definition 2 holds, this generalises ambiguity decompositions stated earlier. For example, when ℓ is the squared loss, it reduces to Krogh and Vedelsby (1994). More generally, it reduces to Theorem 3, the generalised ambiguity decomposition.

The powerful nature of this expression is that we can also consider losses where Definition 2 does *not* hold, e.g., 0/1 or absolute loss. For the 0/1 loss, the centroid is the *mode* of the random variable (Domingos, 2000). For a finite set of ensemble members, this means \bar{q} is a majority vote of the individuals in the ensemble. Or, with the absolute loss, \bar{q} is the *median* prediction of the ensemble members.

Given all of the above, we can now present a novel decomposition for the expected risk of an ensemble using the centroid combiner, isolating the *effect* of diversity.

Theorem 8 (Bias-Variance-Diversity effect decomposition) *Given an ensemble of models $\{\mathbf{q}_i\}_{i=1}^m$ combined by the centroid combiner $\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \left[\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i) \right]$, using any loss ℓ , the expected risk of the ensemble decomposes,*

$$\mathbb{E}_D \left[R(\bar{\mathbf{q}}) \right] = \underbrace{R(\mathbf{y}^*)}_{\text{noise}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \left[R(\mathring{\mathbf{q}}_i) - R(\mathbf{y}^*) \right]}_{\text{average bias-effect}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D \left[R(\mathbf{q}_i) - R(\mathring{\mathbf{q}}_i) \right]}_{\text{average variance-effect}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m \left[R(\mathbf{q}_i) - R(\bar{\mathbf{q}}) \right] \right]}_{\text{diversity-effect}}$$

We obtain four terms: the noise, and the *effects* of average bias, average variance, and diversity. Note that here we are considering each term including the expectation over $P(\mathbf{x}, y)$. For losses where Definition 2 holds, Theorem 8 reduces to Theorem 5, the bias-variance-diversity decomposition.

Similarly to the diversity in Theorem 5, the diversity-*effect* here has the *opposite* sign to the others. Thus, as diversity-effect grows larger, it *reduces* expected risk. However, the diversity-effect can potentially take *negative* values. This may seem non-intuitive, but makes more sense if we consider the term is simply the difference between the average risk and the ensemble risk. So, when the ensemble performs *worse* than the average, this will be negative.

In summary, we have the same conclusion as in the previous sub-section: the effect of diversity can be formulated as a hidden dimension in a decomposition of the ensemble risk. All these terms can be estimated from data—we explore this next.

4.3 Estimating Diversity from Data

We present illustrative experiments, estimating bias/variance/diversity terms from data. We emphasize that we make no claims on the superiority of the centroid combiner over any other, only using it obtain the decomposition. Further experiments in Appendix A.1.

Squared loss: A simple instance of Theorem 5 is squared loss, $\ell(y, q) = (y - q)^2$, which implies $\bar{q} = \frac{1}{m} \sum_{i=1}^m q_i$, and $\hat{q} = \mathbb{E}_D[q]$. The decomposition at each test point (\mathbf{x}, y) is,

$$\mathbb{E}_D \left[(\bar{q} - y)^2 \right] = \underbrace{\frac{1}{m} \sum_{i=1}^m (\hat{q}_i - y)^2}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [(q_i - \hat{q}_i)^2]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m (q_i - \bar{q})^2 \right]}_{\text{diversity}}.$$

Figure 5 shows an experiment with Bagged regression trees, varying ensemble size.




Figure 5: Bagging depth 8 trees, increasing ensemble size (California Housing data).

Whilst expected risk decreases, the average bias and average variance are constant. This is as we might anticipate, since the individual capacities are constant, it is only the ensemble size, m , that we change. In contrast, diversity *increases* with m —subtracting from the expected risk—and the improvement *is determined entirely by diversity*. This is not so if we vary something other than m . Figure 6 fixes $m = 10$, varying tree depth—all *three components now change*, and we have a *bias-variance-diversity* trade-off.




Figure 6: Bagging $m = 10$ regression trees, increasing maximum depth.

Cross-entropy loss: We can also consider Theorem 5 for the KL-divergence between probability vectors, where the centroid combiner $\bar{\mathbf{q}}$ is a normalized geometric mean. For a single (\mathbf{x}, \mathbf{y}) , where \mathbf{y} is a one-hot vector, this gives a decomposition of the cross-entropy, since for any \mathbf{q} , we can write $K(\mathbf{y} \parallel \mathbf{q}) = -\mathbf{y} \cdot \ln \mathbf{q}$, where we use the convention $\mathbf{y} \cdot \ln \mathbf{y} = 0$.

Corollary 9 Take a test point (\mathbf{x}, \mathbf{y}) , where \mathbf{y} is a one-hot vector. Consider an ensemble $\{\mathbf{q}\}_{i=1}^m$ where each model predicts a distribution, combined as $\bar{\mathbf{q}} := Z^{-1} \prod_i \mathbf{q}_i^{1/m}$. We have,

$$\underbrace{-\mathbb{E}_D[\mathbf{y} \cdot \ln \bar{\mathbf{q}}]}_{\text{expected cross-entropy}} = \underbrace{-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \hat{\mathbf{q}}_i}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D[K(\hat{\mathbf{q}}_i \parallel \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D\left[\frac{1}{m} \sum_{i=1}^m K(\bar{\mathbf{q}} \parallel \mathbf{q}_i)\right]}_{\text{diversity}}.$$

With ensembles of neural networks, a normalized geometric mean is equivalent to averaging the network logits, a common practice (Hinton et al., 2015), with both pros and cons (Tassi et al., 2022). It is also known as the ‘product rule’ (Kuncheva, 2014), and has undergone much scrutiny, showing it can perform either better or worse than the arithmetic mean, depending on various factors (Tax et al., 1997). Regardless of performance, the normalized geometric mean is *required* to decompose the cross-entropy in this way. In Section 5 we consider the changes in the decomposition if using the arithmetic mean.

In Figure 7 we compare Bagging MLPs on MNIST, using small networks each with a single layer of 20 hidden nodes, versus larger networks each with 100 hidden nodes. We observe very similar patterns as we saw with squared loss.




Figure 7: Decomposing expected ensemble cross-entropy for Bagging small/larger MLPs.

Overall, the ensemble of larger networks have performed better: this can again be explained by examining the expected risk components. We see in the middle panel that the ensemble of smaller networks have a significantly higher diversity (green line). However, in the far right hand panel, we see the the larger networks have a significantly lower⁵ average bias/variance (blue/orange lines). This gives a significant advantage in the bias/variance/diversity trade-off. Overall, the higher diversity among the small networks is outweighed by the more powerful individuals in the ensemble of larger networks.

5. The lower variance is counter to the ML folklore that increasing model capacity should also increase variance; but this is consistent with recent observations (Yang et al., 2020)

The 0/1 loss and Majority Voting Ensembles: One of the most common ensemble schemes is majority voting, assessed by 0/1 loss (Kuncheva, 2003; Didaci et al., 2013). Here, Theorem 5 does not hold. So, we use the *effect* decomposition of Theorem 8, and the centroid combiner \bar{q} is the majority vote. Using $\mathbb{P}(\bar{q} \neq y) \stackrel{\text{def}}{=} \mathbb{E}_D[\ell_{0/1}(y, \bar{q})]$ for the probability of ensemble error across possible data sets drawn from D , at each test point (\mathbf{x}, y) we have:

$$\mathbb{P}(\bar{q} \neq y) = \underbrace{\frac{1}{m} \sum_{i=1}^m [\ell_{0/1}(\hat{q}_i, y)]}_{\text{average bias-effect}} + \underbrace{\frac{1}{m} \sum_{i=1}^m [\mathbb{P}(q_i \neq y) - \ell_{0/1}(\hat{q}_i, y)]}_{\text{average variance-effect}} - \underbrace{\frac{1}{m} \sum_{i=1}^m [\mathbb{P}(q_i \neq y) - \mathbb{P}(\bar{q} \neq y)]}_{\text{diversity-effect}}.$$

Note that the average bias-effect is simply the proportion of the ensemble members whose centroid \hat{q}_i is incorrect. The diversity-effect is the difference between the average probability of error, and the ensemble probability of error. One might be tempted to re-formulate the diversity-effect, looking for a rearrangement that does not involve the label. In fact, for the 0/1 loss, this is impossible: not just with the majority vote, but with *any* combiner rule.

Theorem 10 (Non-existence of label-independent diversity-effect for 0/1 loss)

For the 0/1 loss, using any ensemble combiner rule, the difference between the average individual risk and the ensemble risk is necessarily dependent on the label.

Therefore, the effect of diversity in majority voting ensembles is necessarily dependent on the label: an unavoidable property of the 0/1 loss. We now compare Bagging and Random Forests on MNIST (10 classes). In the leftmost panel below, we see the Random Forest initially under-performing, but overtaking Bagging as we increase ensemble size.




Figure 8: Bias/Variance/Diversity effect for Bagging vs Random Forests.

In the right hand panels, we decompose the expected risk, and see familiar patterns. As we increase m , the bias/variance-effects are constant, but diversity-effect increases. The variance-effect in Random Forests is much higher than that of Bagging, but this is compensated by higher diversity-effect in larger ensembles. We can also see the average bias-effect is very slightly lower in the Random Forest, which we attribute to the trees being able to avoid noisy features.

Theorem 8 can be extended relatively easily for *weighted* voting—as shown in Appendix C.2. In Figure 9 we plot the ensemble 0/1 risk components for Bagging and two different boosting algorithms, using decision stumps as base learners. Here we adopt a synthetic dataset used by Mease and Wyner (2008) to analyse boosting algorithms.




Figure 9: Mease data, ensembling decision stumps.

We note some interesting differences between the parallel model (Bagging) and the sequential models (LogitBoost/AdaBoost). As before, the bias/variance are constant for Bagging. However for boosting, they vary with m . This is due to the non-homogeneous nature of the individuals, specialising to different parts of the data. With the 0/1 loss, the diversity-effect can be *greater* than the variance-effect, as opposed to the diversity itself, which is upper-bounded by the average variance (see Section 5). More generally, with Bagging the performance comes from increasing diversity. But, the story for boosting is more subtle, with performance determined by a complex trade-off between the components.

4.4 Summary

We presented a framework to understand diversity—explaining it as a measure of model fit, in precisely the same sense as bias/variance, but with the opposite sign. The expected ensemble risk is then determined by a three-way *bias/variance/diversity* trade-off. When an additive bias-variance decomposition is not available, we showed how to instead isolate the *effects* of bias/variance and diversity. In Section 5 we apply this framework with *Bregman divergences* (Bregman, 1967). This will give us a convenient analytical form, and additional insights: notably, a deeper understanding of the centroid combiner rule.

5. Understanding Diversity with Bregman Divergences

We apply our framework to the family of *Bregman divergences* (Bregman, 1967), including many losses as special cases, and enabling us to derive several interesting properties.

5.1 The Basics of Bregman Divergences

A Bregman divergence is defined in terms of a *generator function*, ϕ . Let $\phi : \mathcal{Y} \rightarrow \mathbb{R}$ be a strictly convex function defined on a convex set $\mathcal{Y} \subseteq \mathbb{R}^k$, such that ϕ is differentiable on $\text{ri}(\mathcal{Y})$ —the relative interior of \mathcal{Y} . The Bregman divergence $B_\phi : \mathcal{Y} \times \text{ri}(\mathcal{Y}) \rightarrow \mathbb{R}_+$ is defined,

$$B_\phi(\mathbf{y}, \mathbf{q}) \stackrel{\text{def}}{=} \phi(\mathbf{y}) - \phi(\mathbf{q}) - \langle \nabla \phi(\mathbf{q}), \mathbf{y} - \mathbf{q} \rangle, \quad (14)$$

where $\langle \cdot, \cdot \rangle$ denotes an inner product, and $\nabla \phi(\mathbf{q})$ denotes the gradient vector of ϕ at \mathbf{q} . The choice of ϕ leads to many different losses. With $\phi(q) = q^2$, the gradient vector $\nabla \phi(\mathbf{q})$ is a scalar derivative $d\phi(q)/dq = 2q$, and we recover a squared loss, $B_\phi(y, q) = (y - q)^2$.




Figure 10: Bregman divergence illustrated for the generator $\phi(q) = q^2$. For this example, we have a divergence $B_\phi(y, q) = (y - q)^2 = (1.0 - 0.3)^2 = 0.49$.

Alternatively, we can take a vector $\mathbf{q} \in \mathbb{R}^{k-1}$, for a k -class problem. Note, this is not a probability vector summing to one. It is, however, the minimal description of the distribution, as the k th class probability is $1 - \sum_c \mathbf{q}^{(c)}$. With a particular generator (see final row of Table 2) we recover the the KL-divergence between the distributions in \mathbb{R}^k .

Loss function	$B_\phi(\mathbf{y}, \mathbf{q})$	Generator $\phi(\mathbf{q})$	Domain \mathcal{Y}
Squared loss	$(y - q)^2$	q^2	$q \in \mathbb{R}$
Itakura-Saito	$\frac{y}{q} - \ln \frac{y}{q} - 1$	$-\ln q$	$q \in \mathbb{R}^+$
Poisson loss	$y \ln \frac{y}{q} - (y - q)$	$q \ln q - q$	$q \in \mathbb{R}^+$
KL-divergence	$\sum \mathbf{y}^{(c)} \ln \frac{\mathbf{y}^{(c)}}{\mathbf{q}^{(c)}} + (1 - \sum \mathbf{y}^{(c)}) \ln \frac{1 - \sum \mathbf{y}^{(c)}}{1 - \sum \mathbf{q}^{(c)}}$	$\sum \mathbf{q}^{(c)} \ln \mathbf{q}^{(c)} + (1 - \sum \mathbf{q}^{(c)}) \ln (1 - \sum \mathbf{q}^{(c)})$	$\mathbf{q} \in [0, 1]^{k-1}$ <i>s.t.</i> $\sum_c \mathbf{q}^{(c)} \leq 1$

Table 2: Common loss functions and their Bregman generators.

5.2 Applying our Framework to the family of Bregman Divergences

As discussed in Section 4, we require the existence of a bias-variance decomposition in the form of Definition 2. For Bregman divergences, Pfau (2013) proved the following:

$$\underbrace{\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} [B_\phi(\mathbf{Y}, \mathbf{q})] \right]}_{\text{expected ensemble risk}} = \underbrace{\mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [B_\phi(\mathbf{Y}, \mathbf{Y}^*)] \right]}_{\text{noise}} + \underbrace{B_\phi(\mathbf{Y}^*, \mathring{\mathbf{q}})}_{\text{bias}} + \underbrace{\mathbb{E}_D [B_\phi(\mathring{\mathbf{q}}, \mathbf{q})]}_{\text{variance}}, \quad (15)$$

where the centroid $\mathring{\mathbf{q}}$ is conveniently available in closed-form,

$$\mathring{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D [B_\phi(\mathbf{z}, \mathbf{q})] = [\nabla \phi]^{-1} \left(\mathbb{E}_D [\nabla \phi(\mathbf{q})] \right). \quad (16)$$

This matches Definition 2. If $\phi(q) = q^2$, then $B_\phi(y, q) = (y - q)^2$, and $\mathring{q} = \mathbb{E}_D [q]$, meaning the overall expression is exactly that of Geman et al. (1992). With other generators/losses, the expression corresponds to other bias-variance decompositions, e.g., Heskes (1998). From this one result, using the framework presented in Section 4, we can derive Bregman versions of the ambiguity decomposition, the centroid combiner, and the bias-variance-diversity decomposition.

Theorem 11 (Bregman Ambiguity Decomposition) *For a label $\mathbf{y} \in \mathcal{Y}$ and a set of predictions $\mathbf{q}_1, \dots, \mathbf{q}_m \in \text{ri}(\mathcal{Y})$, combined as $\bar{\mathbf{q}} = [\nabla \phi]^{-1} \left(\frac{1}{m} \sum_{i=1}^m \nabla \phi(\mathbf{q}_i) \right)$. Then we have:*

$$B_\phi(\mathbf{y}, \bar{\mathbf{q}}) = \frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{y}, \mathbf{q}_i) - \frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i). \quad (17)$$

Definition 12 (Bregman Centroid Combiner) *The Bregman centroid combiner is the minimizer of the average divergence from all members. For an ensemble $\{\mathbf{q}_i\}_{i=1}^m$, this is*

$$\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{z}, \mathbf{q}_i) = [\nabla \phi]^{-1} \left(\frac{1}{m} \sum_{i=1}^m \nabla \phi(\mathbf{q}_i) \right). \quad (18)$$

Theorem 13 (Bregman Bias-Variance-Diversity decomposition)

For an ensemble of models $\{\mathbf{q}_i\}_{i=1}^m$, let $\mathring{\mathbf{q}}_i$ be the left Bregman centroid of model \mathbf{q}_i , i.e., $\mathring{\mathbf{q}}_i \stackrel{\text{def}}{=} [\nabla \phi]^{-1} (\mathbb{E}_D [\nabla \phi(\mathbf{q}_i)])$, and define the ensemble output $\bar{\mathbf{q}} \stackrel{\text{def}}{=} [\nabla \phi]^{-1} \left(\frac{1}{m} \sum_{i=1}^m \nabla \phi(\mathbf{q}_i) \right)$. Then we have the decomposition

$$\mathbb{E}_D [\mathbb{E}_{\mathbf{X}\mathbf{Y}} [B_\phi(\mathbf{Y}, \bar{\mathbf{q}})]] = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}} B_\phi(\mathbf{Y}, \mathbf{Y}^*)}_{\text{noise}} + \underbrace{\frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{Y}^*, \mathring{\mathbf{q}}_i)}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [B_\phi(\mathring{\mathbf{q}}_i, \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) \right]}_{\text{diversity}} \right],$$

where $\mathbf{Y}^* = \mathbb{E}_{\mathbf{Y}|\mathbf{X}} [\mathbf{Y}]$.

Examples for different losses (i.e., different Bregman generators) are shown in Table 3. One point that this makes clear is that *the mathematical formulation of diversity is specific to the loss function.*

Expected Ensemble Loss	Average Bias	Average Variance	Diversity
Squared $\mathbb{E}_D [(\bar{q} - y)^2]$	$\frac{1}{m} \sum_{i=1}^m (\hat{q}_i - y)^2$	$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [(q_i - \hat{q}_i)^2]$	$\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m (q_i - \bar{q})^2 \right]$
KL-divergence (Bernoulli) $\mathbb{E}_D \left[y \ln \frac{y}{\bar{q}} + (1-y) \ln \frac{1-y}{1-\bar{q}} \right]$	$\frac{1}{m} \sum_{i=1}^m y \ln \frac{y}{\hat{q}_i} + (1-y) \ln \frac{1-y}{1-\hat{q}_i}$	$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D \left[\hat{q}_i \ln \frac{\hat{q}_i}{q_i} + (1 - \hat{q}_i) \ln \frac{1-\hat{q}_i}{1-q_i} \right]$	$\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m \bar{q} \ln \frac{\bar{q}}{q_i} + (1 - \bar{q}) \ln \frac{1-\bar{q}}{1-q_i} \right]$
KL-divergence (Multinoulli) $\mathbb{E}_D [K(\mathbf{y} \parallel \bar{\mathbf{q}})]$	$\frac{1}{m} \sum_{i=1}^m K(\mathbf{y} \parallel \hat{\mathbf{q}}_i)$	$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [K(\hat{\mathbf{q}}_i \parallel \mathbf{q}_i)]$	$\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m K(\bar{\mathbf{q}} \parallel \mathbf{q}_i) \right]$
Itakura-Saito $\mathbb{E}_D \left[\frac{y}{\bar{q}} - \ln \frac{y}{\bar{q}} - 1 \right]$	$\frac{1}{m} \sum_{i=1}^m \frac{y}{\hat{q}_i} - \ln \frac{y}{\hat{q}_i} - 1$	$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D \left[\frac{\hat{q}_i}{q_i} - \ln \frac{\hat{q}_i}{q_i} - 1 \right]$	$\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m \left(\frac{\bar{q}}{q_i} - \ln \frac{\bar{q}}{q_i} - 1 \right) \right]$
Poisson $\mathbb{E}_D \left[y \ln \frac{y}{\bar{q}} - (y - \bar{q}) \right]$	$\frac{1}{m} \sum_{i=1}^m \left[y \ln \frac{y}{\hat{q}_i} - (y - \hat{q}_i) \right]$	$\frac{1}{m} \sum_{i=1}^m \left[\mathbb{E}_D [q_i] - \hat{q}_i \right]$	$\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m q_i - \prod_{i=1}^m q_i^{1/m} \right]$

Table 3: Bias, variance, and diversity under different Bregman divergences. In all cases, the expectation over $p(\mathbf{x})$ is omitted and the expressions given are for a single point (\mathbf{x}, \mathbf{y}) . See Table 4 for the definitions of the centroid combiners \bar{q} .

5.3 Properties of the Bregman Centroid Combiner

The centroid combiner is the result of a constraint, requiring the ensemble combination $\bar{\mathbf{q}}$ to take the same analytical form as the centroid $\mathring{\mathbf{q}}$, as found in the bias-variance decomposition for the relevant loss. For Bregman divergences, this form (Definition 12) is known in the information geometry literature: as a *left Bregman centroid* (Nielsen and Nock, 2009). In general these turn out to be *quasi-arithmetic* means, and include several well-known ensemble combiner rules, some shown in Table 4.

Loss function	Centroid Combiner	Name
Squared loss	$\frac{1}{m} \sum_{i=1}^m q_i$	Arithmetic mean
Poisson regression loss	$\prod_{i=1}^m q_i^{\frac{1}{m}}$	Geometric mean
KL-divergence	$Z^{-1} \prod_{i=1}^m \left(\mathbf{q}_i^{(c)} \right)^{\frac{1}{m}}$	Normalised geometric mean
Itakura-Saito loss	$1 / \left(\frac{1}{m} \sum_{i=1}^m \frac{1}{q_i} \right)$	Harmonic mean

Table 4: Centroid combiners (i.e., left Bregman centroid of the ensemble) for various losses.

Ensemble averaging in a dual coordinate system: The centroid combiner can be understood as an ensemble averaging operation, *in a new coordinate system*. The mapping between coordinate systems is defined by the gradient of the Bregman generator with respect to its argument, $\nabla \phi(\mathbf{q})$. This is illustrated in Figure 11 for the KL-divergence.




Figure 11: Ensemble averaging in the geometry defined by the KL-divergence.

We use notation q for the *primal* coordinate system, and η for the *dual* coordinate system. In this simple illustration we are predicting a single probability $p \in (0, 1)$. The primal-dual mapping is the gradient $\eta_i = \frac{\partial}{\partial q_i} \phi(q_i) = \ln \frac{q_i}{1-q_i}$, plotted as the blue curve. Two points in the primal $\{q_1 = 0.7, q_2 = 0.97\}$ are mapped to the dual $\{\eta_1 \approx 0.8473, \eta_2 \approx 3.476\}$, then combined via arithmetic mean ($\bar{\eta} \approx 2.16$), and finally mapped back by the inverse operation $\bar{q} = \exp(\bar{\eta}) / (1 + \exp(\bar{\eta})) \approx 0.896$. The centroid combiner is therefore an *arithmetic mean ensemble in the dual coordinate system*, which is equivalent to the left Bregman centroid of the models in the *primal* coordinate system. An equivalent definition was considered by Gupta et al. (2022) under the assumption of i.i.d. models. Our analysis both complements and extends this by removing the i.i.d. assumption, and more fully characterising the properties of ensembles using this combination rule.

Example for KL-divergence of probability estimates: In the case of KL divergence, the centroid combiner is a normalised⁶ geometric mean in the primal coordinate system, i.e., the probability simplex. Figure 12 shows this for the 3-class case.




Figure 12: Combining $m = 4$ predictions in the probability simplex.

The centroid combiner is shown as the blue star (also connected to the 4 individual points) and the *arithmetic* mean as the pink star, which is the centre of mass for the points. For squared loss the resulting Bregman geometry is Euclidean, meaning the centre of mass and the centroid are the same. However, other divergences define a non-Euclidean geometry, meaning the centroid and centre of mass are different. Note that points are connected in the simplex not by straight lines, but by *dual geodesics* defined by the choice of ϕ (Nielsen and Nock, 2009). A different ϕ (and therefore a different loss) would provide a different geometry, and imply a different ensemble combination.

6. Note that the Bregman centroid is only a normalized distribution with ϕ as in Table 2—see Appendix B.4.

5.4 Properties of the Decomposition

We explore further properties of the bias-variance-diversity decomposition.

5.4.1 HOMOGENEOUS VS HETEROGENEOUS ENSEMBLES

An ensemble comprises a set of models $\{\mathbf{q}_i\}_{i=1}^m$, but can also be considered as ‘one big model’ in its own right. Thus, we can discuss the bias/variance of $\bar{\mathbf{q}}$ itself. These bias/variance terms can be related to those of the individual models, using the theorem below.

Theorem 14 *The ensemble bias and ensemble variance can be re-written as a function of the individual models:*

$$\underbrace{B_\phi(\mathbf{y}, \bar{\mathbf{q}})}_{\text{ensemble bias}} = \underbrace{\frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{y}, \hat{\mathbf{q}}_i)}_{\text{average bias}} - \Delta, \quad (19)$$

$$\underbrace{\mathbb{E}_D \left[B_\phi(\bar{\mathbf{q}}, \bar{\mathbf{q}}) \right]}_{\text{ensemble variance}} = \Delta + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D \left[B_\phi(\hat{\mathbf{q}}_i, \mathbf{q}_i) \right]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) \right]}_{\text{diversity}}, \quad (20)$$

where the common term is $\Delta = \frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \hat{\mathbf{q}}_i)$.

The interesting term here is Δ , which we refer to as the model ‘disparity’. This accounts for diversity among the model *families* in the ensemble. If the individual models are all the same family, (e.g., all neural nets of a particular architecture), then the only differences in their predictions will be due to the random variable D . In this case their centroids will be the same $\hat{\mathbf{q}}_i = \hat{\mathbf{q}}_j = \bar{\mathbf{q}}$, and the disparity will be zero. If instead they have different inductive biases, the centroids are not necessarily the same, and the disparity may be non-zero. These are usually referred to as *homogeneous* vs. *heterogeneous* ensembles.

Bias reduction? In a homogeneous ensemble, $\Delta = 0$, and the ensemble bias is equal to the average, i.e., no reduction in bias, relative to the individual models. In a heterogeneous ensemble with sufficiently different model families, $\Delta > 0$, and the bias is *always reduced*.

Variance reduction? Ensembles are often referred to as ‘variance reduction’ methods. We see in Equation (20) that for any Bregman divergence, for a homogeneous ensemble ($\Delta = 0$), the ensemble variance is guaranteed to be less than the average, i.e., we retain the variance reduction property, *provided that we use the centroid combiner rule*. From the same perspective, we see the diversity is upper-bounded by the average variance. The amount by which the ensemble variance is reduced is exactly the diversity of the ensemble. In a heterogeneous ensemble, the story is less simple, as it has both the addition of the disparity, and the subtraction of the diversity.

The properties above apply for *any* Bregman divergence. In practice, we will have a particular choice for a learning problem—we now discuss specific properties of two ubiquitous losses: the cross-entropy, and squared loss.

5.4.2 CROSS ENTROPY LOSS: AVERAGING ESTIMATES OF CLASS PROBABILITIES

When combining class probability estimates, a very popular strategy is to take their arithmetic mean, e.g., Lakshminarayanan et al. (2017), but, if we use the cross-entropy, this is *not* the centroid combiner. We might wonder what effect this has. The proposition below demonstrates that the cross-entropy loss of the ensemble is still *guaranteed to be less than the average loss of its members*, but the ambiguity becomes dependent on the label.

Proposition 15 *Assume a true probability vector, $\mathbf{y} \in \mathbb{R}^k$, and a set of models $\{\mathbf{q}_i\}_{i=1}^m$ combined by averaging, i.e., $\mathbf{q}^\dagger = \frac{1}{m} \sum_{i=1}^m \mathbf{q}_i$, then the cross-entropy loss is*

$$\underbrace{-\mathbf{y} \cdot \ln \mathbf{q}^\dagger}_{\text{ensemble cross-entropy}} = \underbrace{-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \mathbf{q}_i}_{\text{average cross-entropy}} - \underbrace{\sum_{c=1}^k \mathbf{y}^{(c)} \ln \frac{\frac{1}{m} \sum_{j=1}^m \mathbf{q}_j^{(c)}}{\left(\prod_{i=1}^m \mathbf{q}_i^{(c)}\right)^{\frac{1}{m}}}}_{\text{ambiguity (label-dependent)}}, \quad (21)$$

where the second term is non-negative, thus the ensemble loss is guaranteed less than or equal to the average individual loss.

This property can be observed without the framework we have presented thus far, by taking the difference between the $-\mathbf{y} \cdot \ln \mathbf{q}^\dagger$ and $-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \mathbf{q}_i$ terms. In fact, for the case of two classes, this was observed independently by Ivaşcu et al. (2021). Using the general case above, we can combine it with the methodology from Section 4, to obtain a result analogous to the bias-variance-diversity decomposition.

Proposition 16 (Diversity for Averaged Probabilities is label-dependent)

Let $\mathbf{q}^\dagger = \frac{1}{m} \sum_{i=1}^m \mathbf{q}_i$, with $\mathbf{q}_i \in [0, 1]^k$. The expected cross-entropy admits the decomposition:

$$-\mathbb{E}_D [\mathbf{y} \cdot \ln \mathbf{q}^\dagger] = \underbrace{-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \hat{\mathbf{q}}_i}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [K(\hat{\mathbf{q}}_i \parallel \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\sum_{c=1}^k \mathbf{y}^{(c)} \ln \frac{\frac{1}{m} \sum_{j=1}^m \mathbf{q}_j^{(c)}}{\left(\prod_{i=1}^m \mathbf{q}_i^{(c)}\right)^{\frac{1}{m}}} \right]}_{\text{dependency}}.$$

This is very similar to Theorem 9, except that the final term here is dependent on the label \mathbf{y} . To distinguish these, we avoid using the name “diversity”, and instead refer to it as a “dependency” term.

Gupta et al. (2022) also studied properties of ensemble bias/variance with an arithmetic mean combiner, showing that (under an i.i.d. model assumption), the ensemble variance is always reduced. At the same time, they raised a concern, that this may potentially increase the ensemble bias (above the average bias), dependent on the label distribution. Our proposition adds insight: the overall expected loss will *always* be less than the average. Thus, even if there is an increase in ensemble bias, it is always more than compensated by the reduction in ensemble variance, leading to lower overall expected loss.

5.4.3 RELATION TO THE BIAS-VARIANCE-COVARIANCE DECOMPOSITION

For the case of squared loss, our decomposition can be contrasted with the bias-variance-covariance decomposition of Ueda and Nakano (1996), which states, for a point (\mathbf{x}, y) :

$$\mathbb{E}_D [(\bar{q} - y)^2] = \tag{22}$$

$$\underbrace{(\mathbb{E}_D [\bar{q}] - y)^2}_{\text{bias}(\bar{q})} + \underbrace{\frac{1}{m} \frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [(q_i - \mathbb{E}_D [q_i])^2]}_{1/m \times \text{variance}} + \underbrace{\frac{1}{m^2} \sum_{i,j} \mathbb{E}_D [(q_i - \mathbb{E}_D [q_i])(q_j - \mathbb{E}_D [q_j])]}_{(1-1/m) \times \text{covariance}}.$$

Comparing this to our Theorem 13, a difference can be seen in the bias components. Ours is the *average individual* bias, whereas Ueda & Nakano’s is the *ensemble* bias:

$$\overline{\text{bias}} = \frac{1}{m} \sum_{i=1}^m (\mathbb{E}_D [q_i] - y)^2, \tag{23}$$

$$\text{bias}(\bar{q}) = (\mathbb{E}_D [\bar{q}] - y)^2. \tag{24}$$

Ueda & Nakano observed a re-writing of their term: $(\mathbb{E}_D [\bar{q}] - y)^2 = (\frac{1}{m} \sum_{i=1}^m [\mathbb{E}_D [q_i] - y])^2$, and described this as “*the square of the average biases*”. However, we remind the reader that *the square is an artefact of using the squared loss function*. This square is not present in generalised forms of the bias-variance decomposition—see Appendix B.1. Thus, $(\mathbb{E}_D [\bar{q}] - y)^2$ should be referred to as simply the “bias of the ensemble”. The *difference* between $\overline{\text{bias}}$ and $\text{bias}(\bar{q})$ is itself an interesting quantity—if we instantiate (19) for squared loss, we see $\text{bias}(\bar{q}) = \overline{\text{bias}} - \Delta$. This shows that their term is in fact made up of two components: the average of the individual biases, and the *disparity* term introduced above. Of course, if the models are *homogeneous*, i.e., of the same family, then the disparity is zero, and $\text{bias}(\bar{q}) = \overline{\text{bias}}$.

Ueda & Nakano’s decomposition relies on a property of linear combinations of random variables: $\text{Var}(aX_1 + bX_2) = a^2\text{Var}(X_1) + b^2\text{Var}(X_2) + 2ab\text{Cov}(X_1, X_2)$, where X_1, X_2 represent two model outputs and $aX_1 + bX_2$ represents the ensemble combination. When the ensemble combination rule is *non-linear* (i.e., not a simple arithmetic mean) then this property (and hence this as a route to understand diversity) no longer applies.

It is notable that the covariance can be either positive or negative. Our diversity term, however, is always non-negative, growing with more disagreement around the ensemble decision. The covariance is a fundamentally *pairwise* computation—it is likely that this form inspired the many published *pairwise* diversity measures (Kuncheva, 2014). Diversity in the Bregman case is written in a *non-pairwise* manner—the expected average deviation around the ensemble prediction. We conjecture that this term cannot be expressed as solely pairwise operations, implying that pairwise measures may be fundamentally limited.

5.5 Summary

We applied the framework developed in Section 4 to the broad family of Bregman divergences. These have a particularly convenient analytical form, which enabled us to derive several interesting properties.

6. Discussion of Related Work

We provide a discussion of related literature, and outline some possible future work.

6.1 A diversity of diversities

There are *many* proposed measures of ensemble diversity. Kuncheva (2014, Section 8.5) categorises measures along three lines: (1) diversity as a characteristic of the classifiers, not considering the combiner or the true label; (2) diversity as a characteristic of the classifiers and the combiner, but *not* using the true label; and (3) diversity as a characteristic of the classifiers, the combiner, *and* the true label. Our framework clearly supports the second idea, for losses where Definition 2 holds (e.g., Bregman divergences), and the third, for losses where it does not (e.g., 0/1 loss, absolute loss).

6.2 Good and Bad Diversity in Majority Voting Ensembles

With the 0/1 loss, Theorem 5 does not apply. As a consequence, we have to use the *effect* decomposition, in Theorem 8. As a reminder, we still consider the *measurement* of diversity in the same form as other losses: $\mathbb{E}_D[\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(\bar{q}, q_i)]$. However, the *effect* that this quantity has on the ensemble risk is given by the *diversity-effect* term. This can be related to the idea of ‘good’ and ‘bad’ diversity (Brown and Kuncheva, 2010).

“Good” vs “Bad” diversity. Brown and Kuncheva (2010) showed that, restricting the label to $y \in \{-1, +1\}$, and \bar{q} as a majority vote, the following holds:

$$\underbrace{\ell_{0/1}(y, \bar{q})}_{\text{ensemble loss}} = \underbrace{\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i)}_{\text{average loss}} - \underbrace{y\bar{q} \frac{1}{m} \sum_{i=1}^m \ell_{0/1}(\bar{q}, q_i)}_{\text{good/bad diversity}}. \quad (25)$$

The similarity to the ambiguity decompositions from earlier is self-evident. The difference is the $y\bar{q}$ preceding the second term. **The sign of $y\bar{q}$ mediates the effect of the diversity.** If $y\bar{q} = +1$ (equivalent to saying $\bar{q} = y$), the term is non-negative and therefore acts to subtract from the average error. When this occurs, it is referred to as “good” diversity. When $y\bar{q} = -1$, the diversity adds to the error, referred to as “bad” diversity. Comparing this to the Ambiguity-effect decomposition (Theorem 7) we can relate the two as follows.

Corollary 17 *For 2-class problems, $y \in \{-1, +1\}$, the ambiguity-effect is Brown & Kuncheva’s good/bad diversity term, averaged over the data distribution.*

$$\underbrace{\mathbb{E}_{\mathbf{X}Y} \left[Y\bar{q} \frac{1}{m} \sum_{i=1}^m \ell_{0/1}(\bar{q}, q_i) \right]}_{\text{good/bad diversity}} = \underbrace{\frac{1}{m} \sum_{i=1}^m \left[R_{0/1}(q_i) - R_{0/1}(\bar{q}) \right]}_{\text{ambiguity-effect}}. \quad (26)$$

The concepts of Good/Bad diversity were generalised to the multi-class case by Didaci et al. (2013), and the same relation applies. The idea was used in Bian and Chen (2019) to define diversity in weighted ensembles, giving the same form as Didaci et al. (2013).

When the effect of diversity is negative. The definition of diversity-effect is simple to state: *the expected difference between the average individual risk and the ensemble risk*. If the ensemble risk performs worse than the individuals, this will be negative, and hence the effect of diversity will be to *increase* the 0/1 risk. Assume a majority voting ensemble of classifiers, making independent errors on a k -class problem, with each model predicting the correct class with probability p . If $p > 0.5$, then the ensemble 0/1 risk is guaranteed to be less than the individual 0/1 risk (Lam and Suen, 1997). In our context this means the diversity-effect will be non-negative, and thus the effect of diversity is to reduce ensemble risk. However, this idealised scenario will certainly not always be possible. Kuncheva (2014, Section 8.3) details carefully constructed scenarios where more diversity is associated with *increasing* ensemble risk, i.e., the bad diversity outweighs the good diversity.

6.3 The Correlation of Diversity and Classification Accuracy

Figure 2 is a toy “accuracy/diversity” scatter plot, in the style popularised by Kuncheva and Whitaker (2003); it shows the correlation of a diversity measure to accuracy (0/1 loss). The x-axis is some diversity measure, and the y-axis is $\frac{1}{m} \sum_{i=1}^m R_{0/1}(q_i) - R_{0/1}(\bar{q})$, i.e., the difference between the average individual error and the ensemble error. The higher this value, the more the ensemble outperforms the average individual model. A higher correlation to diversity is seen to be a more successful diversity measure, as it explains the performance improvement. Figure 13 shows two such plots for *real* data, for the cross-entropy diversity. We estimate diversity on validation data, and the 0/1 loss on a final test set. The caveat here, is that we have assumed access to class probability estimates, for computing the cross-entropy diversity term. This may not always be possible.

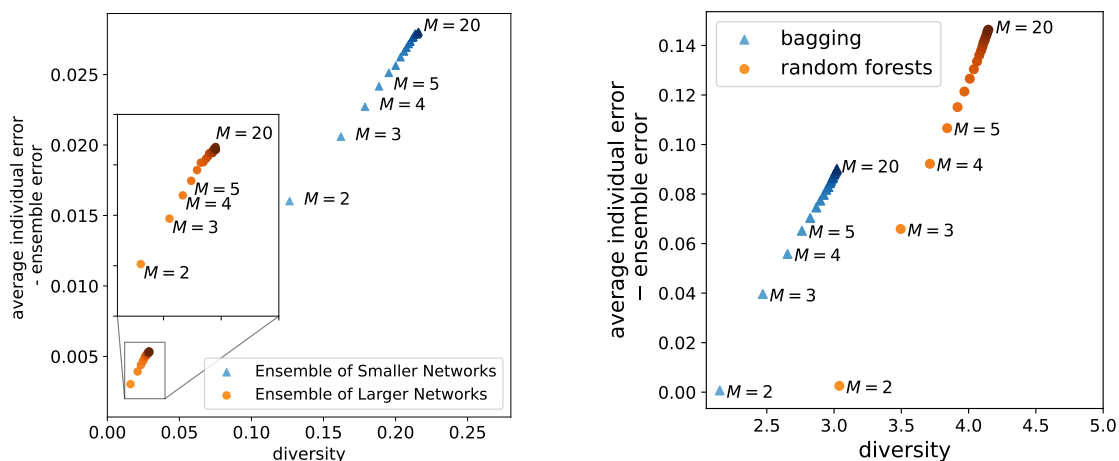


Figure 13: Accuracy/Diversity plots. LEFT: Bagged MLPs (corresponds to experiment in Figure 7). RIGHT: Decision tree ensembles, Bagging vs Random Forests.

We see a strong correlation for both ensembles. The smaller networks (blue triangles, Pearson’s $r^2 = 0.998$) have greater diversity than the larger networks (orange circles, $r^2 = 0.992$). The plot must be read in the context that overall, the larger networks significantly

outperformed the smaller networks—it simply shows that the performance came from more powerful base models, as opposed to their diversity.

In the right hand figure, we compare Bagging decision trees (unlimited depth) and Random Forests on MNIST. In both cases the trees are combined by obtaining probabilities and combining via normalised geometric mean. Again we see strong correlation of diversity and performance gain. Bagging has a correlation $r^2 \approx 0.996$, whilst Random Forests has $r^2 \approx 0.999$. For a fixed m , we can compare corresponding points, where the *only difference is the additional split-point randomisation of the forest*. At $m = 20$, RF provides a reduction in generalisation error of $\approx 14.5\%$, versus only $\approx 9\%$ for Bagging. It interesting to note this is solely due to increased diversity generated by random feature splits.

One might wonder *why* there is such a strong correlation in both cases. If we remember the alternative view of the same experiment, Figure 7, we see that bias/variance are constant, and it is *only the diversity that changes*. When any change is observed in the overall ensemble cross-entropy, we know it is caused by a change in diversity. Therefore, if we can assume strong correlation between the ensemble cross-entropy and the 0/1 loss, then there will be a similar strong correlation between diversity and 0/1 loss.

We might now wonder, with this diversity measure, *will we always see a strong correlation between diversity and reduction in 0/1 loss?* The answer is no, for a very good reason that highlights a critical point in our understanding of diversity. In Figure 14 we fix at $m = 10$ bagged trees, and vary their depth. The expected loss reduces—however, now it is not solely due to diversity. Now, the bias and variance also change rapidly, and the correlation of 0/1 loss/diversity is much lower.




Figure 14: Bagging $m = 10$ trees, varying depth, and correlation is now $r^2 \approx 0.59$.

When varying any other parameter than m , one *should not expect* to see a strong correlation of performance improvement and diversity. This is because, if we vary any parameter that alters individual capacities, then the average bias/variance also changes, and diversity is not the only factor in play. The overall performance is decided by a 3-way trade-off, just as there is a 2-way trade-off of bias/variance in single models. It would be interesting to explore this with ensembles of very deep neural networks, where the bias/variance trade-off seems to not act as classical theory predicts (Belkin et al., 2018).

6.4 Future work

Re-visiting previous work involving heuristic diversity measures. In the introduction, we outlined several areas that had made use of heuristic diversity measures. Examples were given for diversity helping defend adversarial attacks, or counteracting covariate shift, or in forming more computationally efficient models. Any one of these might be repeated, using diversity derived from the loss function, as in our framework.

Forcing diversity. A natural line of research might be to *enforce diversity* in some sense, i.e., using our diversity measures as regularisers during the construction of an ensemble. Negative Correlation (NC) Learning (Liu and Yao, 1999) encourages diversity in regression ensembles. Brown et al. (2006) showed this is in fact exploiting the squared loss ambiguity decomposition, Equation (1). The Bregman ambiguity decomposition, Equation (17), implies that the NC algorithm is a special case of a wider family of diversity-encouraging losses—the case for cross-entropy was explored in Webb et al. (2021).

Diversity for Margin losses. Margin losses (e.g., logistic or exponential loss) are an important class of losses, best known in the context of *Gradient Boosting* algorithms. Wood et al. (2022) analysed bias-variance decompositions for such losses. Using these results, it would be possible to obtain bias-variance-diversity decompositions for margin losses, and thus potentially gain insight into boosting algorithms. However, Mease and Wyner (2008) showed strong evidence that the additive model form in AdaBoost/LogitBoost results in a disconnect between the surrogate margin loss and our true objective, the 0/1 loss. In particular, the surrogate loss can go *up* (sometimes exponentially fast) whilst the 0/1 loss on a hold-out sample is going *down*. This implies that any analysis of the surrogate (including loss decompositions) does not necessarily give meaningful insights for the 0/1 loss. Furthermore, with boosting, the individual models are more naturally interpreted as learning to correct the errors of previous ensemble members rather than perform well in their own right, making interpretation of the average bias term problematic. Extending our framework for these cases could be an interesting new interpretation of this popular class of algorithms.

Diversity in Mixtures of Experts. As it stands, our framework does not apply for a mixture of experts model, where the models are gated by an additional input-dependent model, itself learnt from data. Future work might consider diversity for this scenario—perhaps measuring diversity of predictions in small regions of the space, where the gated weights of members are reasonably constant.

Understanding or Encoding Probabilistic Assumptions. The concept of diversity is one way of measuring dependencies between predictions. An explicit *probabilistic* assumption of conditional independence (sometimes known as the *Dawid-Skene* model) would seemingly imply some degree of diversity. It is not clear how the loss function view of diversity (which we present) relates to this probabilistic view. Given the duality between losses/log-likelihood, and that Bregman divergences can be seen as KL-divergences of probability densities, this should be possible and may yield interesting insights.

7. Conclusion

We have presented a unified theory of ensemble diversity. A key insight is that it is not the task (e.g., classification/regression) that matters, but the *loss function*. We demonstrated that diversity can be seen as a *hidden dimension in the bias-variance decomposition of an ensemble loss*. Diversity emerges naturally with this point of view—the exact functional form is specific to the loss being used, but the decompositions have a common structure applicable for a wide range of losses:

$$\text{expected loss} = (\text{average bias}) + (\text{average variance}) - (\text{diversity}).$$

This gives a clear relationship between the ensemble performance and its diversity. The only other scenario where this was previously available is for squared loss with an arithmetic mean combiner (Ueda and Nakano, 1996). Our framework is an alternative in this case, but generalises the notion of diversity to a far wider range of losses, including the cross-entropy, and Poisson regression losses, but more generally any Bregman divergence. The decomposition requires the use of a particular combination rule, specific to the loss at hand, which we call the *centroid combiner*. This combiner turns out to correspond to several well-known combiners already in the literature—e.g., for cross-entropy, it is the normalized geometric mean. This generalises the idea of ensemble “averaging” to many other scenarios, explained as averaging in a dual coordinate system defined by a Bregman divergence.

For losses where an additive bias-variance decomposition does not exist, we adopted the approach of James and Hastie (1997), to instead measure the *effects* of bias, variance, and diversity, which turn out to be dependent on the label distribution. The case of 0/1 loss is particularly interesting—we show that, not just for majority voting, but for any combiner, the effect of diversity is necessarily a *label-dependent* quantity.

We therefore have a broad and precise formulation of diversity for a wide range of supervised learning scenarios. This challenge has been referred to as the “holy grail” of ensemble learning (Zhou, 2012, pg 100), an open question for over 30 years. The answer we provide phrases diversity as a measure of *model fit*, in precisely the same sense as bias/variance. Thus, we should not be aiming to “maximise diversity” as so many works aim to do. Instead, just as bias and variance change with model characteristics, the same applies to diversity, and we have to manage the three-way *bias/variance/diversity trade-off*.

Code for all experiments at: <https://github.com/EchoStatements/Decompose>

Acknowledgments

Funding in direct support of this work: EPSRC EP/N035127/1 (LAMBDA project) and EP/T026995/1 (EnnCore project). Mikel Luján holds an Arm/RAEng Research Chair Award and a Royal Society Wolfson Fellowship.

GB would like to thank LK and FR for a career’s worth of inspiration & support.

Appendix A. Additional Experimental details

A.1 Additional results

We present additional results on several datasets. We emphasize that we make no claims on empirical superiority of any one method over any other. We simply use these toy datasets as illustrative examples of how the risk components can be estimated.

Squared loss: We use California housing data—Table 5 shows results from three ensembles (each $M = 30$ regression trees), compared to a single tree. We use Bagging with constrained trees (max depth 8) and compare against unlimited depth trees, and a Random Forest.

Single tree (depth 8)	Bagging (depth 8)	Bagging (unconstrained)	Random Forest
0.47	0.35	0.30	0.28

Table 5: California housing data: MSE of a single tree versus ensembles of 30 trees.

We observe that the Random Forest is the best choice here, followed up closely by the unconstrained Bagging. Figure 15 explains their performance by decomposing risk into bias, variance, and diversity—also showing how the components change as we grow the ensemble.




Figure 15: Decomposing the expected risk of three decision tree ensembles.

We observe the same behaviour as in Figure 5. The diversity increases with m , and is upper-bounded by the value of the average variance. A higher average variance effectively raises the “ceiling” to which diversity can rise. The average variance is higher as we move from depth-limited trees to unlimited depth, and higher again with the random split-points in the Random Forest (here we use the square root of the number of features). The higher average variance is compensated for by the diversity, causing Random Forest to be the best option. It is notable that for large ensembles, the expected risk of the ensemble is almost entirely due to the value of the average bias (≈ 0.28 in the case of unconstrained trees), with diversity having essentially cancelled out the average variance of the individual models. This behaviour is not just a quirk of this data set, in fact it holds as long as the individuals are all from the same model family, i.e., the ensemble is *homogeneous*—the general case is discussed in Section 5.4.

Cross-entropy: In Figure 16, we have additional accuracy/diversity plots for neural network ensembles on classification problems (n examples, d features, k outputs/classes). Datasets: Phoneme ($n = 5404$, $d = 4$, $k = 2$), Landsat ($n = 6435$, $d = 36$, $k = 6$), Spambase ($n = 4601$, $d = 57$, $k = 2$), South German Credit (2019 version: $n = 1000$, $d = 20$, $k = 2$).

In each case, the squared Pearson’s correlation coefficient is shown in the legend. The following configuration was used in all MLP experiments:

- learning rate: 0.1 (Stochastic gradient descent)
- num epochs: 50 (MNIST), 200 (other data sets)
- hidden layer size (20 small/100 larger)
- number of trials: 100

where each trial uses a 90% sub-sample of the full training data, as outlined in Figure 17.




Figure 16: Error/diversity plots across four data sets, comparing ensembles of small MLPs (20 hidden nodes, blue dots) versus large (100 hidden nodes, orange dots).

A.2 Methodology for Estimating Bias, Variance, and Diversity

We present our methodology for estimating the bias, variance and diversity from data. Algorithm 1 shows the procedure for experiments where we estimate diversity of ensembles of different sizes, such as in the experiments for Bagging and Random Forests. Notably, an ensemble of size $m + 1$ is created by using the members of the ensemble of size m , rather creating a new ensemble of size $m + 1$ from scratch. We also present a visualisation of the sub-sampling scheme used for Bagging in Figure 17.

m : number of ensemble members
 n : number of test points
 t : number of trials
 D_k : the training set to use in trial k

Arguments: $model, train_data, test_data$

Output: $test_preds$: an array of model predictions of size $t \times m \times n$.

```

for  $k \in \{1, \dots, t\}$  do
  for  $j \in \{1, \dots, n\}$  do
     $D_k \leftarrow$  90% of  $train\_data$ , sampled without replacement;
    for  $i \in \{1, \dots, m\}$  do
       $member\_data \leftarrow$  bootstrap from  $D_k$ ;
       $i$ th ensemble member  $\leftarrow$  copy of  $model$  trained on  $member\_data$ ;
       $test\_preds[k, i, j] \leftarrow$  prediction of  $test\_data$  from the  $i$ th ensemble
        member, in the  $k$ th trial,  $j$ th test data point;
    end
  end
end

```

Algorithm 1: Algorithm for collecting data, later used to estimate diversity of a Bagging ensemble, whilst varying ensemble size m .

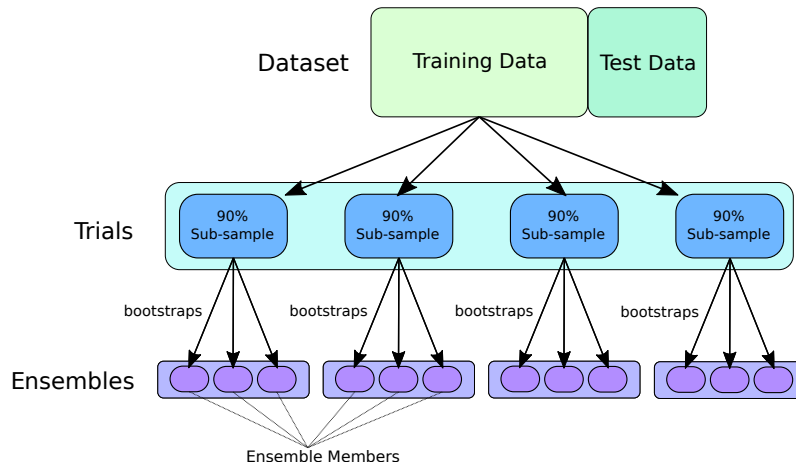


Figure 17: Visualisation of the sub-sampling scheme used for Bagging ensembles.

The result of Algorithm 1 is an array of size (t, m, n) . For Bregman divergences, the average bias and average variance are approximated by first estimating the left Bregman centroid for each ensemble member, replacing the expectation over the random variable D with an average over t trials. We use notation D_k as the full training data in the k th trial, from which the Bagging algorithm can be applied. The centroid for model i on a single test point \mathbf{x} is therefore estimated as,

$$\hat{q}_i^{est}(\mathbf{x}) = [\nabla\phi]^{-1} \left(\frac{1}{t} \sum_{k=1}^t \nabla\phi(q_i(\mathbf{x}; D_k)) \right) \approx \hat{q}_i(\mathbf{x})$$

With this estimate of $\hat{q}_i(\mathbf{x})$, the average bias and average variance are computed as

$$\text{average bias} \approx \frac{1}{m} \frac{1}{n} \sum_{i=1}^m \sum_{j=1}^n B_\phi(y_j, \hat{q}_i^{est}(\mathbf{x}_j)), \quad (27)$$

$$\text{average variance} \approx \frac{1}{t} \frac{1}{m} \frac{1}{n} \sum_{k=1}^t \sum_{i=1}^m \sum_{j=1}^n B_\phi(\hat{q}_i^{est}(\mathbf{x}_j), q_i(\mathbf{x}_j; D_k)). \quad (28)$$

Diversity is estimated similarly, with \bar{q} defined as the centroid combiner rule applied to the m ensemble members in a given trial:

$$\text{diversity} \approx \frac{1}{t} \frac{1}{m} \frac{1}{n} \sum_{k=1}^t \sum_{i=1}^m \sum_{j=1}^n B_\phi(\bar{q}(\mathbf{x}_j, D_k), q_i(\mathbf{x}_j; D_k)). \quad (29)$$

Appendix B. Further explanations

We present further exposition on certain points within the paper. The material in this appendix is not novel, but expands on selected aspects to support our argument.

B.1 Is it “squared bias”, or just “bias”?

Geman et al. (1992) presented a squared loss bias-variance decomposition, Equation (3). One of the terms is $(\mathbb{E}_D [q] - y^*)^2$. They, and many subsequent papers, refer to this as “squared bias”, or “bias squared”. *This is a misnomer*, as we will now argue.

The root of the issue is that Geman et al. borrowed terminology from classical statistics. For some population value θ , and an estimator $\hat{\theta}$, classical statistics refers to the quantity $\mathbb{E}[\hat{\theta}] - \theta$ as the *bias* of $\hat{\theta}$. Thus, by analogy, they refer to $(\mathbb{E}_D [q] - y^*)^2$ as the *squared bias* of estimator q , where y^* is assumed to be the population value.

The reason this is a misnomer—is that similar decompositions are known for other losses, and the square is *not present*, e.g. Heskes (1998); Pfau (2013); Wood et al. (2022). Relevant to our paper, the ensemble bias-variance-covariance decomposition of Ueda and Nakano (1996) unfortunately also inherited the “squared” misnomer, as we explained in sub-section 5.4.3. These results were proven long after the publication of Geman et al. (1992), thus we cannot blame the authors, especially since the term “bias” is overloaded in science (e.g. inductive, sampling, selection, confirmation). *So, in short, the term “squared bias” is a misnomer, and throughout our paper we refer to just “bias”.*

B.2 The Ambiguity Decomposition is equivalent to Bias-Variance

In the bias-variance decomposition, the random variable D is commonly assumed to be over all possible training sets of a fixed size n . This is a convention introduced by Geman et al. (1992). However this is not necessarily the case, and D can be over any stochastic quantity involved in construction of the model (e.g. random initial weights for neural networks), or indeed over a finite set of *pre-constructed models*. In this latter case, we see an equivalence between the bias-variance decomposition and the ambiguity decomposition—in the sense that if one exists, then the other must also. This can be seen most easily by considering a simpler form of the bias-variance decomposition, with no noise:

$$\mathbb{E}_D [\ell(y, q)] = \ell(y, \hat{q}) + \mathbb{E}_D [\ell(\hat{q}, q)]. \quad (30)$$

The relation still holds with noise, but is easier to explain with this assumption. Now, define D as a discrete random variable with an event space consisting of m different pre-constructed models, $\{q_i\}_{i=1}^m$, and the probability assigned to each event is uniform, i.e., $1/m$ for each event q_i . This means the expectation \mathbb{E}_D becomes an average $\frac{1}{m} \sum_{i=1}^m$. The centroid is still defined in the same way, but we use notation \bar{q} to denote the fact that it concerns a finite set of predictions: $\bar{q} \stackrel{\text{def}}{=} \arg \min_{z \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m \ell(z, q_i)$. This means the expression above becomes:

$$\frac{1}{m} \sum_{i=1}^m \ell(y, q_i) = \ell(y, \bar{q}) + \frac{1}{m} \sum_{i=1}^m \ell(\bar{q}, q_i). \quad (31)$$

A simple rearrangement of terms yields the ambiguity decomposition.

B.3 Properties of Heskes' Decomposition for the Cross-Entropy

Heskes (1998) presents a bias-variance decomposition for the KL-divergence between probability densities, which he extends to expose a noise term, giving a noise-bias-variance decomposition for the cross-entropy. For Multinoulli distributions, we claim this is equivalent to that we present in Equation (4), which is important as it then matches our unified bias-variance form. For clarity, we restate our KL decomposition here:

$$\underbrace{\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} [K(\mathbf{Y} \parallel \mathbf{q})] \right]}_{\text{expected risk}} = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [K(\mathbf{Y} \parallel \mathbf{Y}^*)]}_{\text{noise}} + \underbrace{K(\mathbf{Y}^* \parallel \mathring{\mathbf{q}})}_{\text{bias}} + \underbrace{\mathbb{E}_D [K(\mathring{\mathbf{q}} \parallel \mathbf{q})]}_{\text{variance}} \right], \quad (32)$$

where \mathbf{Y} is random variable, over one-hot encoded indicators of the true class for an input point \mathbf{x} . If we take Heskes' expression (Heskes, 1998, Eq 6) and assume Multinoulli distributions, we get a decomposition of the expected cross-entropy at a single point \mathbf{x} , between a true class distribution \mathbf{y}^* and a model \mathbf{q} . In our notation this reads:

$$-\mathbb{E}_D [\mathbf{y}^* \cdot \ln \mathbf{q}] = \underbrace{-\mathbf{y}^* \cdot \ln \mathbf{y}^*}_{\text{noise}} + \underbrace{K(\mathbf{y}^* \parallel \mathring{\mathbf{q}})}_{\text{bias}} + \underbrace{\mathbb{E}_D [K(\mathring{\mathbf{q}} \parallel \mathbf{q})]}_{\text{variance}}, \quad (33)$$

where $\mathring{\mathbf{q}} \propto \exp(\mathbb{E}_D[\ln \mathbf{q}])$ is the normalized geometric mean of the model distribution.

To show the equivalence, we first note that (32) has an expectation over $P(\mathbf{X})$. We can restate this at a single \mathbf{x} ,

$$\underbrace{\mathbb{E}_D \left[\mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} [K(\mathbf{Y} \parallel \mathbf{q})] \right]}_{\text{expected loss}} = \underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} [K(\mathbf{Y} \parallel \mathbf{y}^*)]}_{\text{noise}} + \underbrace{K(\mathbf{y}^* \parallel \mathring{\mathbf{q}})}_{\text{bias}} + \underbrace{\mathbb{E}_D [K(\mathring{\mathbf{q}} \parallel \mathbf{q})]}_{\text{variance}}, \quad (34)$$

where $\mathbf{y}^* = \mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}}[\mathbf{Y}]$ is the true class distribution at \mathbf{x} . The second and third terms on the right are equal to those in Heskes' decomposition. For the noise term we note,

$$\mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} [K(\mathbf{Y} \parallel \mathbf{y}^*)] = \mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} \left[\mathbf{Y} \cdot \ln \frac{\mathbf{Y}}{\mathbf{y}^*} \right] \quad (35)$$

$$= \mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} [\mathbf{Y} \cdot \ln \mathbf{Y}] - \mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} [\mathbf{Y} \cdot \ln \mathbf{y}^*] \quad (36)$$

$$= -\mathbb{E}_{\mathbf{Y}|\mathbf{X}=\mathbf{x}} [\mathbf{Y}] \cdot \ln \mathbf{y}^* \quad (37)$$

$$= -\mathbf{y}^* \cdot \ln \mathbf{y}^*, \quad (38)$$

where we used the convention $0 \ln 0 = 0$. Since the right hand sides are equal, therefore the left hand sides are equivalent too.

B.4 The Importance of Parameter Encoding in the KL-divergence

Consider the KL-divergence between two discrete probability vectors of length k ,

$$K(\mathbf{p} \parallel \mathbf{q}) = \sum_{c=1}^k \mathbf{p}^{(c)} \ln \frac{\mathbf{p}^{(c)}}{\mathbf{q}^{(c)}} \quad (39)$$

There are two ways in which this can be expressed as a Bregman divergence, either using the full-length probability vectors, $\mathbf{p} \in \mathbb{R}^k$ and using the generator

$$\phi_{\text{full}}(\mathbf{p}) = \sum_{c=1}^k \mathbf{p}^{(c)} \ln \mathbf{p}^{(c)} \quad (40)$$

or using the minimally parameterised vectors, $\tilde{\mathbf{p}} \in \mathbb{R}^{k-1}$, where the last entry is omitted:

$$\phi_{\text{min}}(\tilde{\mathbf{p}}) = \sum_{c=1}^{k-1} \tilde{\mathbf{p}}^{(c)} \ln \tilde{\mathbf{p}}^{(c)} + (1 - \sum_{c'=1}^{k-1} \tilde{\mathbf{p}}^{(c')}) \ln(1 - \sum_{c'=1}^{k-1} \tilde{\mathbf{p}}^{(c')}). \quad (41)$$

Given two probability vectors in the appropriate form, either formulation gives a Bregman divergence is equivalent to the KL-divergence between the vectors of class probability estimates (Nielsen and Nock, 2009), i.e.,

$$B_{\phi_{\text{full}}}(\mathbf{p}, \mathbf{q}) = B_{\phi_{\text{min}}}(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}) = K(\mathbf{p} \parallel \mathbf{q}).$$

We use the minimal parameterisation, as it exhibits a desirable property. In particular, it is necessary to use the second to ensure that the Bregman centroid is always a valid distribution on the probability simplex. In this case, the centroid combiner is the normalised geometric mean, as we now demonstrate.

Minimally parameterized vectors: To show this we consider m minimally parameterised vectors, $\mathbf{q}_i \in \mathbb{R}^{k-1}$ (note that we have dropped tilde above \mathbf{q} for simplicity). Our claim is that the centroid combiner is the normalised geometric mean $\bar{\mathbf{q}} = [\nabla \phi]^{-1} \left(\frac{1}{m} \sum_{i=1}^m \nabla \phi(\mathbf{q}_i) \right)$, is of the form

$$\bar{\mathbf{q}}^{(c)} = [\nabla \phi_{\text{min}}]^{-1} \left(\frac{1}{m} \sum_{i=1}^m \nabla \phi_{\text{min}}(\mathbf{q}_i) \right) = \frac{\prod_{i=1}^m \mathbf{q}_i^{(c) \frac{1}{m}}}{\sum_{c'=1}^k \prod_{i=1}^m \hat{\mathbf{q}}^{(c') \frac{1}{m}}}, \quad (42)$$

where $\hat{\mathbf{q}}$ denotes the extension of the $k-1$ length vector into a full k length probability vector. Plugging in the gradients from Table 6, we start with

$$\bar{\mathbf{q}}^{(c)} = \frac{\exp \left(\frac{1}{m} \sum_{i=1}^m \ln \frac{\mathbf{q}_i^{(c)}}{1 - \sum_{c'=1}^{k-1} \mathbf{q}_i^{(c')}} \right)}{1 + \sum_{c'=1}^{k-1} \exp \left(\frac{1}{m} \sum_{i=1}^m \ln \frac{\mathbf{q}_i^{(c')}}{1 - \sum_{c''=1}^{k-1} \mathbf{q}_i^{(c'')}} \right)}.$$

Note that the numerator here can be rearranged:

$$\exp \left(\frac{1}{m} \sum_{i=1}^m \ln \frac{\mathbf{q}_i^{(c)}}{1 - \sum_{c'=1}^{k-1} \mathbf{q}_i^{(c')}} \right) = \prod_{i=1}^m \left(1 - \sum_{c'=1}^{k-1} \mathbf{q}_i^{(c')} \right)^{-\frac{1}{m}} \prod_{i=1}^m \left(\mathbf{q}_i^{(c)} \right)^{\frac{1}{m}},$$

and the denominator can be written

$$\begin{aligned}
1 + \sum_{c'=1}^{k-1} \exp \left(\frac{1}{m} \sum_{i=1}^m \ln \frac{\mathbf{q}_m^{(c')}}{1 - \sum_{c''=1}^{k-1} \mathbf{q}_m^{(c'')}} \right) &= 1 + \sum_{c'=1}^{k-1} \prod_{i=1}^m \mathbf{q}_i^{(c') \frac{1}{m}} \left(1 - \sum_{c''=1}^{k-1} \mathbf{q}_i^{(c'')} \right)^{-\frac{1}{m}} \\
&= 1 + \prod_{i'=1}^m \left(1 - \sum_{c''=1}^{k-1} \mathbf{q}_{i'}^{(c'')} \right)^{-\frac{1}{m}} \sum_{c'=1}^{k-1} \prod_{i=1}^m \mathbf{q}_i^{(c') \frac{1}{m}} \\
&= \prod_{i'=1}^m \left(1 - \sum_{c''=1}^{k-1} \mathbf{q}_{i'}^{(c'')} \right)^{-\frac{1}{m}} \left(\prod_{i=1}^m \left(1 - \sum_{c''=1}^{k-1} \mathbf{q}_i^{(c'')} \right)^{\frac{1}{m}} + \sum_{c'=1}^{k-1} \prod_{i=1}^m \mathbf{q}_i^{(c') \frac{1}{m}} \right).
\end{aligned}$$

Putting the numerator and denominator back into the second expression of Equation (42) and using the definition of $\hat{\mathbf{q}}$, we find the first terms in both products cancel and we are left with the required result.

Full length k Probability Vector: If we do not use the minimally parameterized vectors, we would have the Bregman generator $\phi(\mathbf{p}) = \sum_{c=1}^k \mathbf{p}^{(c)} \ln \mathbf{p}^{(c)}$. This gives the geometric mean, rather than the normalised version. To see this, we first note that

$$\begin{aligned}
(\nabla \phi(\mathbf{p}))^{(c)} &= 1 + \ln \mathbf{p}^{(c)} = \boldsymbol{\eta}^{(c)} \\
([\nabla \phi]^{-1}(\boldsymbol{\eta}))^{(c)} &= \exp(\boldsymbol{\eta}^{(c)} - 1),
\end{aligned}$$

and therefore the centroid combiner is

$$\begin{aligned}
\left([\nabla \phi]^{-1} \left(\frac{1}{m} \sum_{m=1}^m \nabla \phi(\mathbf{q}_m) \right) \right)^{(c)} &= \exp \left(\frac{1}{m} \sum_{i=1}^m 1 + \ln \mathbf{q}_i^{(c)} - 1 \right) \\
&= \exp \left(\frac{1}{m} \sum_{i=1}^m \ln \mathbf{q}_i^{(c)} \right) = \prod_{i=1}^m \mathbf{q}_i^{(c) \frac{1}{m}}.
\end{aligned}$$

Note that this means that $\bar{\mathbf{q}}$ is not necessarily a valid probability vector. In fact, it is a valid probability vector *only if* $\mathbf{q}_1 = \dots = \mathbf{q}_m$.

B.5 List of Bregman Centroids

For quick reference purposes (see Nielsen and Nock (2009) for more details) we list centroids for different losses of interest to the ML community. The left Bregman centroid is defined:

$$\mathring{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D [B_\phi(\mathbf{z}, \mathbf{q})] = [\nabla \phi]^{-1} \left(\mathbb{E}_D [\nabla \phi(\mathbf{q})] \right). \quad (43)$$

The centroid $\mathring{\mathbf{q}}$ takes different forms dependent on the generator used.

Loss	Gradient $\eta = \nabla \phi(\mathbf{q})$	Inverse Grad. $\mathbf{q} = [\nabla \phi]^{-1}(\eta)$	Left Bregman Centroid $\mathring{\mathbf{q}} := [\nabla \phi]^{-1} \left(\mathbb{E}_D [\nabla \phi(\mathbf{q})] \right)$
Squared	$2q$	$\frac{1}{2}\eta$	$\mathbb{E}_D[q]$
Itakura-Saito	$-\frac{1}{q}$	$-\frac{1}{\eta}$	$1 / \left(\mathbb{E}_D [1/q] \right)$
Poisson loss	$\ln q$	$\exp(\eta)$	$\exp(\mathbb{E}_D [\ln q])$
KL-divergence	$\ln \frac{\mathbf{q}^{(c)}}{1 - \sum_{c'=1}^{k-1} \mathbf{q}^{(c')}}$	$\frac{\exp(\boldsymbol{\eta}^{(c)})}{1 + \sum_{c'=1}^{k-1} \exp(\boldsymbol{\eta}^{(c')})}$	$\frac{1}{Z} \exp \left(\mathbb{E}_D [\ln \mathbf{q}] \right)$

Table 6: Common losses (see Table 2) and their Bregman centroids. In the case of KL-divergence, Z is a normalizer to ensure a valid distribution.

Appendix C. Proofs for Section 4

C.1 Proofs for Bias-Variance-Diversity and Effect Decompositions

Theorem 5 (Generalized Bias-Variance-Diversity decomposition) *Consider a set of models $\{\mathbf{q}_i\}_{i=1}^m$, evaluated by a loss ℓ . Assuming a bias-variance decomposition holds in the form of Definition 2, the following decomposition also holds.*

$$\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} [\ell(\mathbf{Y}, \bar{\mathbf{q}})] \right] = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [\ell(\mathbf{Y}, \mathbf{Y}^*)]}_{\text{noise}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{Y}^*, \hat{\mathbf{q}}_i)}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [\ell(\hat{\mathbf{q}}_i, \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m \ell(\bar{\mathbf{q}}, \mathbf{q}_i) \right]}_{\text{diversity}} \right], \quad (9)$$

where $\hat{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \mathbb{E}_D [\ell(\mathbf{z}, \mathbf{q})]$ and the combiner is $\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i)$.

Proof Take the expected risk of $\bar{\mathbf{q}}$, and apply Proposition 3, the generalized ambiguity decomposition:

$$\mathbb{E}_D [\mathbb{E}_{\mathbf{X}\mathbf{Y}} [\ell(\mathbf{Y}, \bar{\mathbf{q}})]] = \mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} \left[\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{Y}, \mathbf{q}_i) \right] \right] - \mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} \left[\frac{1}{m} \sum_{i=1}^m \ell(\bar{\mathbf{q}}, \mathbf{q}_i) \right] \right]. \quad (44)$$

Now apply Definition 2, the generalised bias-variance decomposition, to the first term on the right:

$$\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} \left[\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{Y}, \mathbf{q}_i) \right] \right] = \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [\ell(\mathbf{Y}, \mathbf{Y}^*)] + \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{Y}^*, \hat{\mathbf{q}}_i) + \frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [\ell(\hat{\mathbf{q}}_i, \mathbf{q}_i)] \right]. \quad (45)$$

Plugging Equation (45) into (44) completes the proof. \blacksquare

Further explanation on dependent/independent training schemes: For the bias-variance decomposition of a single model q , we defined a random variable $D \sim P(\mathbf{x}, y)^n$, i.e. over i.i.d. training sets of size n . When we have an ensemble of models, it is common for each to have their own training set, e.g. the Bagging algorithm. In this case, D is redefined to a vector $[D_0, D_1, D_2, \dots, D_m]$, where $D_0 \sim P(\mathbf{x}, y)^n$ defines the overall training set supplied to the ensemble. Each D_i , for $1..m$, is a random variable defining an i.i.d. sample with replacement from the same fixed training set defined by D_0 . Thus each D_i is dependent on D_0 , but not on any other D_j . However, it can also be that D_t is dependent on D_{t-1} , for example as in Boosting algorithms. In either case, if we consider any function of (without loss of generality) the first model, q_1 , then by the law of total expectation, $\mathbb{E}_Y [\mathbb{E}_{X|Y}[X]] = \mathbb{E}_X[X]$, we have:

$$\mathbb{E}_{D_0 \dots D_m} [q_1(D_1)] = \mathbb{E}_{D_0, D_2 \dots D_m} \left[\mathbb{E}_{D_1|D_0, D_2 \dots D_m} [q_1(D_1)] \right] = \mathbb{E}_{D_1} [q_1(D_1)]. \quad (46)$$

Thus, the function of the model is only dependent on the individual model's training set.

Theorem 8 (Bias-Variance-Diversity effect decomposition) *Given an ensemble of models $\{\mathbf{q}_i\}_{i=1}^m$ combined by the centroid combiner $\bar{\mathbf{q}} \stackrel{\text{def}}{=} \arg \min_{\mathbf{z} \in \mathcal{Y}} \left[\frac{1}{m} \sum_{i=1}^m \ell(\mathbf{z}, \mathbf{q}_i) \right]$, using any loss ℓ , the expected risk of the ensemble decomposes,*

$$\mathbb{E}_D \left[R(\bar{\mathbf{q}}) \right] = \underbrace{R(\mathbf{y}^*)}_{\text{noise}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \left[R(\mathring{\mathbf{q}}_i) - R(\mathbf{y}^*) \right]}_{\text{average bias-effect}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D \left[R(\mathbf{q}_i) - R(\mathring{\mathbf{q}}_i) \right]}_{\text{average variance-effect}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m \left[R(\mathbf{q}_i) - R(\bar{\mathbf{q}}) \right] \right]}_{\text{diversity-effect}}$$

Proof Note that several terms on the right cancel, reducing to the left-hand side. \blacksquare

Theorem 10 (Non-existence of label-independent diversity-effect for 0/1 loss)

For the 0/1 loss, using any ensemble combiner rule, the difference between the average individual risk and the ensemble risk is necessarily dependent on the label.

Proof We show that $\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q})$ is necessarily dependent on the value of y at a point (\mathbf{x}, y) . Taking expectation over $P(\mathbf{x}, y)$ proves the result. We show a case, where there is no ensemble combination rule such that the expression is independent of y .

The two-class case: Define $\mathcal{Y} = \{1, 2\}$. For a fixed \mathbf{x} , without loss of generality, let $p = 0.6$ be the proportion of the m models predicting class 1. The combiner \bar{q} , provides a label $\in \mathcal{Y}$. For both possible \bar{q} , the quantity $\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q})$ depends on y . First, if we assume $\bar{q} = 1$, we have:

$$\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q}) = \begin{cases} 0.4 - 0 = 0.4 & \text{if } y = 1 \\ 0.6 - 1 = -0.4 & \text{if } y = 2 \end{cases}$$

Alternatively, when $\bar{q} = 2$,

$$\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q}) = \begin{cases} 0.4 - 1 = -0.6 & \text{if } y = 1 \\ 0.6 - 0 = 0.6 & \text{if } y = 2 \end{cases}$$

For both \bar{q} , the value of $\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q})$ is dependent on the true label y .

The multiclass case: Define $\mathcal{Y} = \{1, 2, \dots, k\}$, and the proportion of the m models predicting each class as p_1, \dots, p_k . We set $p_1 = 0.6, p_2 = 0.4$, and zero for all other classes. From the two-class case we know that when $\bar{q} \in \{1, 2\}$, there is a dependency on y . This persists for $\bar{q} \in \{3, \dots, k\}$, where we have

$$\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q}) = \begin{cases} 0.4 - 1 = -0.6 & \text{if } y = 1 \\ 0.6 - 1 = -0.4 & \text{if } y = 2 \\ 1.0 - 0 = 1 & \text{if } y = \bar{q} \\ 1.0 - 1 = 0 & \text{for all other classes.} \end{cases}$$

For all $\bar{q} \in \mathcal{Y}$, the expression $\frac{1}{m} \sum_{i=1}^m \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q})$ is dependent on the label y . \blacksquare

C.2 Bias/Variance/Diversity-Effect Decomposition: Weighted Voting

We now prove a bias-variance-diversity-effect decomposition for a weighted majority vote combination. Given an ensemble of classification models q_1, \dots, q_m each outputting a label prediction from $\{1, \dots, k\}$, and weights for those models $\alpha_1(D), \dots, \alpha_m(D)$, we consider the ambiguity-effect decomposition for weighted plurality vote. The weighted majority vote is

$$\bar{q}(\mathbf{x}; D) = \arg \min_{c \in \{1, \dots, k\}} \sum_{i=1}^m \frac{\alpha_i(D)}{\sum_{j=1}^m \alpha_j(D)} \ell_{0/1}(c, q_i),$$

similarly, the centroid for an ensemble member is given by

$$\mathring{q}_i = \arg \min_{c \in \{1, \dots, k\}} \mathbb{E}_D \left[\frac{\alpha_i(D)}{\mathbb{E}_D[\alpha_i(D)]} \ell_{0/1}(c, q_i) \right],$$

with ties broken randomly (the tie break procedure can be thought of as part of the random variable D , since D implicitly contains all sources of stochasticity related to the model).

Theorem 18 (Ambiguity-Effect Decomposition for Weighted Majority Vote)

With this, we can define a weighted effect decomposition as

$$\ell_{0/1}(y, \bar{q}) = \underbrace{\sum_{i=1}^m a_i(D) \ell_{0/1}(y, q_i)}_{\text{weighted average loss}} - \underbrace{\left[\sum_{i=1}^m a_i(D) \ell_{0/1}(y, q_i) - \ell_{0/1}(y, \bar{q}) \right]}_{\text{ambiguity-effect}},$$

where $a_i = \frac{\alpha_i}{\sum_{j=1}^m \alpha_j}$.

The validity of this theorem can be verified simply by cancelling terms on the right-hand side. The theorem tells us that the loss of an ensemble can be decomposed into a non-negative term (the weighted average loss of the ensemble members), and an ambiguity-effect term, which quantifies how much better (or worse) the performance of the ensemble is compared to the average member loss. Using the same principle, we can also construct a bias/variance-effect decomposition which takes into account a weighting $\alpha(D)$.

Theorem 19 (Bias/variance-effect Decomposition for Weighted Majority Vote)

$$\mathbb{E}_D [\alpha \ell_{0/1}(y, q)] = \mathbb{E}_D [\alpha] \ell_{0/1}(y, \mathring{q}) + \left[\mathbb{E}_D [\alpha \ell_{0/1}(y, q)] - \mathbb{E}_D [\alpha] \ell_{0/1}(y, \mathring{q}) \right]$$

Again, the proof of the result is immediate from considering which terms on the right cancel. However, it is worth considering how the decomposition works and what the terms mean. Consider the decomposition when we replace α with a normalised version $\frac{\alpha}{\mathbb{E}_D[\alpha]}$, we get:

$$\mathbb{E}_D \left[\frac{\alpha}{\mathbb{E}_D[\alpha]} \ell_{0/1}(y, q) \right] = \ell_{0/1}(y, \mathring{q}) + \left[\mathbb{E}_D \left[\frac{\alpha}{\mathbb{E}_D[\alpha]} \ell_{0/1}(y, q) \right] - \ell_{0/1}(y, \mathring{q}) \right].$$

This is exactly bias-variance-effect decomposition that we have seen previously, but re-weighting the contributions of the different data sets. In fact, the two are equivalent, with the weights defining a new probability density function. Taking $P_D(\mathcal{D})$ as the probability density function over data sets, the decomposition above is exactly the bias-variance-effect decomposition with the new probability density function $Q_D(\mathcal{D}) = P_D(\mathcal{D}) \frac{\alpha(\mathcal{D})}{\mathbb{E}_D[\alpha(\mathcal{D})]}$. We can also easily reintroduce label noise and expose a noise term, and turning the bias into a bias-effect:

$$\mathbb{E}_Y \left[\mathbb{E}_D \left[\frac{\alpha}{\mathbb{E}_D[\alpha]} \ell_{0/1}(Y, q) \right] \right] = \underbrace{\mathbb{E}_Y [\ell_{0/1}(Y, Y^*)]}_{\text{noise}} + \underbrace{\mathbb{E}_Y [\ell_{0/1}(Y, \hat{q}) - \ell_{0/1}(Y, Y^*)]}_{\text{weighted bias-effect}} \\ \underbrace{\left[\mathbb{E}_D \left[\mathbb{E}_Y \left[\frac{\alpha}{\mathbb{E}_D[\alpha]} \ell_{0/1}(Y, f) - \ell_{0/1}(Y, \hat{q}) \right] \right] \right]}_{\text{weighted variance-effect}}.$$

We can now apply similar proof methodology, and get the following bias-variance-diversity effect decomposition for weighted majority vote.

Proposition 20 (Bias-Variance-Diversity-Effect for Weighted Voting) *Given m classifiers q_1, \dots, q_m , where the ensemble is a weighted majority vote, i.e., $\bar{q} = \arg \min_{z \in \mathcal{Y}} \sum_{i=1}^m \alpha_i \ell_{0/1}(z, q_i)$ for weights $\alpha_1, \dots, \alpha_m \in \mathbb{R}_+$, the ensemble loss admits the following decomposition, where the normalised weight is $a_i = \alpha_i / \sum_{j=1}^m \alpha_j$.*

$$\mathbb{E}_Y [\mathbb{E}_D [\ell_{0/1}(Y, \bar{q})]] = \underbrace{\mathbb{E}_Y [\ell_{0/1}(Y, Y^*)]}_{\text{noise}} + \underbrace{\sum_{i=1}^m \mathbb{E}_D [a_i \mathbb{E}_Y [\ell_{0/1}(Y, \hat{q}_i) - \ell_{0/1}(Y, Y^*)]]}_{\text{weighted average bias-effect}} \\ + \underbrace{\sum_{i=1}^m \mathbb{E}_D [a_i \mathbb{E}_Y [\ell_{0/1}(Y, q_i) - \ell_{0/1}(Y, \hat{q}_i)]]}_{\text{weighted average variance-effect}} - \underbrace{\mathbb{E}_D \left[\mathbb{E}_Y \left[\sum_{i=1}^m a_i \ell_{0/1}(Y, q_i) - \ell_{0/1}(Y, \bar{q}) \right] \right]}_{\text{diversity-effect}},$$

where $\hat{q}_i = \arg \min_{z \in \mathcal{Y}} \mathbb{E}_D [\alpha_i \ell_{0/1}(z, q_i)]$, noting that α_i and q_i are both dependent on D .

As before, the veracity of this result can be seen by cancelling terms on the right-hand side. AdaBoost produces a set of binary classifiers $h_i \in \{-1, +1\}$ and corresponding weights $\alpha_i(D) \in \mathbb{R}$, so setting $q_i = h_i$ allows immediate application of the decomposition. LogitBoost does not produce classifier/weight pairs, but instead a set of regression models each $g_i \in \mathbb{R}$. We can apply the decomposition by separating these into sign/magnitude components, giving a classification $q_i = \text{sign}(g_i)$ and weight: $\alpha_i(D) = |g_i|$.

Appendix D. Proofs for Section 5

D.1 Proof of Bregman Ambiguity / Bias-Variance-Diversity decompositions.

Technically, the first two theorems in this sub-section are a corollary to the existence of a bias-variance decomposition for Bregman divergences (Pfau, 2013) combined with Theorem 5. However, for completeness and didactic purposes, we present separate proofs here. The first proof will use the Bregman three-point property (Banerjee et al., 2005):

Definition 21 (Bregman three-point property) *Given a convex set $\mathcal{S} \subseteq \mathbb{R}^k$, a strictly convex function $\phi : \mathcal{Y} \rightarrow \mathbb{R}$, and any $\mathbf{x}_1 \in \mathcal{Y}$ and $\mathbf{x}_2, \mathbf{x}_3 \in \text{ri}(\mathcal{S})$, the following identity holds:*

$$B_\phi(\mathbf{x}_1, \mathbf{x}_3) = B_\phi(\mathbf{x}_1, \mathbf{x}_2) + B_\phi(\mathbf{x}_2, \mathbf{x}_3) + \langle \mathbf{x}_1 - \mathbf{x}_2, \nabla\phi(\mathbf{x}_2) - \nabla\phi(\mathbf{x}_3) \rangle. \quad (47)$$

Theorem 11 (Bregman Ambiguity Decomposition) *For a label $\mathbf{y} \in \mathcal{Y}$ and a set of predictions $\mathbf{q}_1, \dots, \mathbf{q}_m \in \text{ri}(\mathcal{Y})$, combined as $\bar{\mathbf{q}} = [\nabla\phi]^{-1}(\frac{1}{m} \sum_{i=1}^m \nabla\phi(\mathbf{q}_i))$. Then we have:*

$$B_\phi(\mathbf{y}, \bar{\mathbf{q}}) = \frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{y}, \mathbf{q}_i) - \frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i). \quad (17)$$

Proof We instantiate the three-point property with $\mathbf{x}_1 = \mathbf{y}$, $\mathbf{x}_2 = \bar{\mathbf{q}}$, and $\mathbf{x}_3 = \mathbf{q}_i$.

$$B_\phi(\mathbf{y}, \mathbf{q}_i) = B_\phi(\mathbf{y}, \bar{\mathbf{q}}) + B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) + \langle \mathbf{y} - \bar{\mathbf{q}}, \nabla\phi(\bar{\mathbf{q}}) - \nabla\phi(\mathbf{q}_i) \rangle \quad (48)$$

Now average both sides of this expression over the m models:

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{y}, \mathbf{q}_i) &= B_\phi(\mathbf{y}, \bar{\mathbf{q}}) + \frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) + \frac{1}{m} \sum_{i=1}^m \left[\langle \mathbf{y} - \bar{\mathbf{q}}, \nabla\phi(\bar{\mathbf{q}}) - \nabla\phi(\mathbf{q}_i) \rangle \right] \\ &= B_\phi(\mathbf{y}, \bar{\mathbf{q}}) + \frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) + \langle \mathbf{y} - \bar{\mathbf{q}}, \nabla\phi(\bar{\mathbf{q}}) - \frac{1}{m} \sum_{i=1}^m \nabla\phi(\mathbf{q}_i) \rangle \\ &= B_\phi(\mathbf{y}, \bar{\mathbf{q}}) + \frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i), \end{aligned} \quad (49)$$

where in the final step we have rearranged the definition of $\bar{\mathbf{q}}$ to use that $\nabla\phi(\bar{\mathbf{q}}) = \frac{1}{m} \sum_{i=1}^m \nabla\phi(\mathbf{q}_i)$. Finally, rearranging the terms of (49) recovers the desired result, Equation (17). \blacksquare

Theorem 13 (Bregman Bias-Variance-Diversity decomposition)

For an ensemble of models $\{\mathbf{q}_i\}_{i=1}^m$, let $\mathring{\mathbf{q}}_i$ be the left Bregman centroid of model \mathbf{q}_i , i.e., $\mathring{\mathbf{q}}_i \stackrel{\text{def}}{=} [\nabla\phi]^{-1}(\mathbb{E}_D[\nabla\phi(\mathbf{q}_i)])$, and define the ensemble output $\bar{\mathbf{q}} \stackrel{\text{def}}{=} [\nabla\phi]^{-1}(\frac{1}{m}\sum_{i=1}^m \nabla\phi(\mathbf{q}_i))$. Then we have the decomposition

$$\mathbb{E}_D[\mathbb{E}_{\mathbf{X}\mathbf{Y}}[B_\phi(\mathbf{Y}, \bar{\mathbf{q}})]] = \mathbb{E}_{\mathbf{X}} \left[\underbrace{\mathbb{E}_{\mathbf{Y}|\mathbf{X}} B_\phi(\mathbf{Y}, \mathbf{Y}^*)}_{\text{noise}} + \underbrace{\frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{Y}^*, \mathring{\mathbf{q}}_i)}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D[B_\phi(\mathring{\mathbf{q}}_i, \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D\left[\frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i)\right]}_{\text{diversity}} \right],$$

where $\mathbf{Y}^* = \mathbb{E}_{\mathbf{Y}|\mathbf{X}}[\mathbf{Y}]$.

Proof Take the expected risk of $\bar{\mathbf{q}}$, and apply the Bregman ambiguity decomposition:

$$\mathbb{E}_D[\mathbb{E}_{\mathbf{X}\mathbf{Y}}[B_\phi(\mathbf{Y}, \bar{\mathbf{q}})]] = \mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} \left[\frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{Y}, \mathbf{q}_i) \right] \right] - \mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} \left[\frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) \right] \right]. \quad (50)$$

Now apply Pfau's decomposition, Equation (15), to the first term on the RHS, and we have

$$\mathbb{E}_D \left[\mathbb{E}_{\mathbf{X}\mathbf{Y}} \left[\frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{Y}, \mathbf{q}_i) \right] \right] = \mathbb{E}_{\mathbf{X}} \left[\mathbb{E}_{\mathbf{Y}|\mathbf{X}} [B_\phi(\mathbf{Y}, \mathbf{Y}^*)] + \frac{1}{m} \sum_{i=1}^m B_\phi(\mathbf{Y}^*, \mathring{\mathbf{q}}_i) + \frac{1}{m} \sum_{i=1}^m \mathbb{E}_D [B_\phi(\mathring{\mathbf{q}}_i, \mathbf{q}_i)] \right]. \quad (51)$$

Plugging Equation (51) into (50) completes the proof. \blacksquare

Theorem 14 The ensemble bias and ensemble variance can be re-written as a function of the individual models:

$$\underbrace{B_\phi(\mathbf{y}, \mathring{\mathbf{q}})}_{\text{ensemble bias}} = \frac{1}{m} \sum_{i=1}^m \underbrace{B_\phi(\mathbf{y}, \mathring{\mathbf{q}}_i)}_{\text{average bias}} - \Delta, \quad (19)$$

$$\underbrace{\mathbb{E}_D [B_\phi(\mathring{\mathbf{q}}, \bar{\mathbf{q}})]}_{\text{ensemble variance}} = \Delta + \frac{1}{m} \sum_{i=1}^m \underbrace{\mathbb{E}_D [B_\phi(\mathring{\mathbf{q}}_i, \mathbf{q}_i)]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\frac{1}{m} \sum_{i=1}^m B_\phi(\bar{\mathbf{q}}, \mathbf{q}_i) \right]}_{\text{diversity}}, \quad (20)$$

where the common term is $\Delta = \frac{1}{m} \sum_{i=1}^m B_\phi(\mathring{\mathbf{q}}, \mathring{\mathbf{q}}_i)$.

Proof

Equation (19) can be proven by applying Theorem 11 to a set of centroid models, $\{\mathring{\mathbf{q}}_i\}_{i=1}^m$. Equation (20) can be proven by applying Theorem 13 but substituting $\mathbf{y} = \mathring{\mathbf{q}}$. \blacksquare

D.2 Proofs for Section 5.4

Proposition 15 Assume a true probability vector, $\mathbf{y} \in \mathbb{R}^k$, and a set of models $\{\mathbf{q}_i\}_{i=1}^m$ combined by averaging, i.e., $\mathbf{q}^\dagger = \frac{1}{m} \sum_{i=1}^m \mathbf{q}_i$, then the cross-entropy loss is

$$\underbrace{-\mathbf{y} \cdot \ln \mathbf{q}^\dagger}_{\text{ensemble cross-entropy}} = \underbrace{-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \mathbf{q}_i}_{\text{average cross-entropy}} - \underbrace{\sum_{c=1}^k \mathbf{y}^{(c)} \ln \frac{\frac{1}{m} \sum_{j=1}^m \mathbf{q}_j^{(c)}}{\left(\prod_{i=1}^m \mathbf{q}_i^{(c)}\right)^{\frac{1}{m}}}}_{\text{ambiguity (label-dependent)}}, \quad (52)$$

where the second term is non-negative, thus the ensemble loss is guaranteed less than or equal to the average individual loss.

Proof Take the average cross-entropy, and subtract the ensemble cross entropy:

$$\begin{aligned} -\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \mathbf{q}_i - \left[-\mathbf{y} \cdot \ln \mathbf{q}^\dagger \right] &= \sum_{c=1}^k \mathbf{y}^{(c)} \ln \mathbf{q}^{\dagger(c)} - \frac{1}{m} \sum_{i=1}^m \sum_{c=1}^k \mathbf{y}^{(c)} \ln \mathbf{q}_i^{(c)} \\ &= \sum_{c=1}^k \mathbf{y}^{(c)} \ln \mathbf{q}^{\dagger(c)} - \sum_{c=1}^k \mathbf{y}^{(c)} \ln \left(\prod_i \mathbf{q}_i^{(c)} \right)^{1/m} \\ &= \sum_{c=1}^k \mathbf{y}^{(c)} \ln \left(\frac{\mathbf{q}^{\dagger(c)}}{\left(\prod_{i=1}^m \mathbf{q}_i^{(c)} \right)^{\frac{1}{m}}} \right) \end{aligned}$$

Using the definition of \mathbf{q}^\dagger and rearranging completes the derivation. From the arithmetic-geometric mean inequality, $\mathbf{q}^{\dagger(c)} \geq \left(\prod_{i=1}^m \mathbf{q}_i^{(c)} \right)^{1/m}$, implying that the term inside the logarithm is greater or equal to 1, and so the overall term is non-negative. \blacksquare

Proposition 16 Let $\mathbf{q}^\dagger = \frac{1}{m} \sum_{i=1}^m \mathbf{q}_i$, with $\mathbf{q}_i \in [0, 1]^k$. The expected cross-entropy admits the decomposition:

$$\begin{aligned} -\mathbb{E}_D \left[\mathbf{y} \cdot \ln \mathbf{q}^\dagger \right] &= \\ &\underbrace{-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \hat{\mathbf{q}}_i}_{\text{average bias}} + \underbrace{\frac{1}{m} \sum_{i=1}^m \mathbb{E}_D \left[K(\hat{\mathbf{q}}_i \parallel \mathbf{q}_i) \right]}_{\text{average variance}} - \underbrace{\mathbb{E}_D \left[\sum_{c=1}^k \mathbf{y}^{(c)} \ln \frac{\frac{1}{m} \sum_{j=1}^m \mathbf{q}_j^{(c)}}{\left(\prod_{i=1}^m \mathbf{q}_i^{(c)} \right)^{\frac{1}{m}}} \right]}_{\text{dependency}}. \end{aligned}$$

Proof Starting with Equation (52) and taking the expectation over D , we have

$$-\mathbb{E}_D \left[\mathbf{y} \cdot \ln \mathbf{q}^\dagger \right] = \mathbb{E}_D \left[-\frac{1}{m} \sum_{i=1}^m \mathbf{y} \cdot \ln \mathbf{q}_i \right] - \mathbb{E}_D \left[\sum_{c=1}^k \mathbf{y}^{(c)} \ln \frac{\frac{1}{m} \sum_{j=1}^m \mathbf{q}_j^{(c)}}{\left(\prod_{i=1}^m \mathbf{q}_i^{(c)} \right)^{\frac{1}{m}}} \right].$$

Now, apply the KL bias-variance decomposition to the first term, and the result is proven. \blacksquare

References

- Kartik Audhkhasi, Abhinav Sethy, Bhuvana Ramabhadran, and Shrikanth Narayanan. Generalized ambiguity decomposition for understanding ensemble diversity. *arXiv:1312.7463*, 2013.
- Arindam Banerjee, Srujana Merugu, Inderjit S Dhillon, and Joydeep Ghosh. Clustering with bregman divergences. *Journal of machine learning research*, 6(Oct):1705–1749, 2005.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine learning and the bias-variance trade-off. *arXiv preprint arXiv:1812.11118*, 2018.
- Yijun Bian and Huanhuan Chen. When does diversity help generalization in classification ensembles? *arXiv preprint arXiv:1910.13631*, 2019.
- Battista Biggio, Iginio Corona, Giorgio Fumera, Giorgio Giacinto, and Fabio Roli. Bagging classifiers for fighting poisoning attacks in adversarial classification tasks. In *Intl workshop on multiple classifier systems*, pages 350–359. Springer, 2011.
- Lev M Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967.
- Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Gavin Brown. An information theoretic perspective on multiple classifier systems. In *International Workshop on Multiple Classifier Systems*, pages 344–353. Springer, 2009.
- Gavin Brown and Ludmila Kuncheva. “good” and “bad” diversity in majority vote ensembles. In *Workshop on multiple classifier systems*, pages 124–133. Springer, 2010.
- Gavin Brown, Jeremy Wyatt, Rachel Harris, and Xin Yao. Diversity creation methods: a survey and categorisation. *Information Fusion*, 6(1):5 – 20, 2005.
- Gavin Brown, Jeremy L Wyatt, and Peter Tiño. Managing diversity in regression ensembles. *Journal of machine learning research*, 6(Sep):1621–1650, 2006.
- Sebastian Buschjäger, Lukas Pfahler, and Katharina Morik. Generalized negative correlation learning for deep ensembling. *arXiv preprint arXiv:2011.02952*, 2020.
- Yue Cao, Thomas Andrew Geddes, Jean Yee Hwa Yang, and Pengyi Yang. Ensemble deep learning in bioinformatics. *Nature Machine Intelligence*, pages 1–9, 2020.
- Miguel A Carreira-Perpinán and Ramin Raziperchikolaei. An ensemble diversity approach to supervised binary hashing. In *Neural Information Processing Systems*, 2016.
- Luca Didaci, Giorgio Fumera, and Fabio Roli. Diversity in classifier ensembles: Fertile concept or dead end? In *Workshop on multiple classifier systems*. Springer, 2013.

- Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000.
- Pedro Domingos. A unified bias-variance decomposition. In *International Conference on Machine Learning*, pages 231–238, 2000.
- Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.
- Pierre Geurts. *Contributions to decision tree induction: bias/variance tradeoff and time series classification*. PhD thesis, University of Liège Belgium, 2002.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- Neha Gupta, Jamie Smith, Ben Adlam, and Zelda E Mariet. Ensembles of classifiers: a bias-variance perspective. *Transactions of Machine Learning Research*, 2022. URL <https://openreview.net/forum?id=1I0QFVncY9>.
- Lars K. Hansen and Peter Salamon. Neural network ensembles. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12:993 – 1001, 11 1990. doi: 10.1109/34.58871.
- Tom Heskes. Bias/variance decompositions for likelihood-based estimators. *Neural Computation*, 10(6):1425–1433, 1998.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Carina Ivaşcu, Richard Everson, and Jonathan Fieldsend. Optimising diversity in classifier ensembles of classification trees. In *Intl Conf. on Applications of Evolutionary Computation.*, 2021.
- Gareth James and Trevor Hastie. Generalizations of the bias/variance decomposition for prediction error. Technical report, Dept. Statistics, Stanford Univ., Stanford, CA, 1997.
- Muhammad Zohaib Jan and Brijesh Verma. A novel diversity measure and classifier selection approach for generating ensemble classifiers. *IEEE Access*, 7, 2019.
- Zhengshen Jiang, Hongzhi Liu, Bin Fu, and Zhonghai Wu. Generalized ambiguity decompositions for classification with applications in active learning and unsupervised ensemble pruning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Ron Kohavi, David H Wolpert, et al. Bias plus variance decomposition for zero-one loss functions. In *ICML*, volume 96, pages 275–83, 1996.
- Dan Kondratyuk, Mingxing Tan, Matthew Brown, and Boqing Gong. When ensembling smaller models is more efficient than single large models. *arXiv preprint arXiv:2005.00570*, 2020.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. Imagenet classification with deep convolutional neural networks. In *Neural information processing systems*, 2012.

- Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation and active learning. In *Neural Information Processing Systems*, pages 231–238, 1994.
- Ludmila Kuncheva. That elusive diversity in classifier ensembles. In *Iberian conference on pattern recognition and image analysis*, pages 1126–1138. Springer, 2003.
- Ludmila Kuncheva. *Combining pattern classifiers: methods and algorithms*. 2nd Edition, John Wiley & Sons, 2014.
- Ludmila Kuncheva and Christopher Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine learning*, 51(2), 2003.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Neural Information Processing Systems*, 2017.
- Louisa Lam and SY Suen. Application of majority voting to pattern recognition: an analysis of its behavior and performance. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 27(5):553–568, 1997.
- Yong Liu and Xin Yao. Ensemble learning via negative correlation. *Neural networks*, 12(10):1399–1404, 1999.
- David Mease and Abraham Wyner. Evidence contrary to the statistical view of boosting. *Journal of Machine Learning Research*, 9(Feb):131–156, 2008.
- Frank Nielsen and Richard Nock. Sided and symmetrized bregman centroids. *IEEE transactions on Information Theory*, 55(6):2882–2904, 2009.
- David W Opitz and Jude W Shavlik. Generating accurate and diverse members of a neural-network ensemble. In *Neural information processing systems*, pages 535–541, 1996.
- Luis A Ortega, Rafael Cabañas, and Andrés R Masegosa. Diversity and generalization in neural network ensembles. *AISTATS 2022, arXiv preprint arXiv:2110.13786*, 2022.
- Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. *International Conference on Machine Learning*, 2019.
- Michael P Perrone and Leon N Cooper. When networks disagree: Ensemble methods for hybrid neural networks. Technical report, Brown University, Institute for Brain and Neural Systems, 1992.
- David Pfau. A Generalized Bias-Variance Decomposition for Bregman Divergences. Technical report, Columbia University, 2013.
- Robert Schapire, Yoav Freund, Peter Bartlett, and Wee Sun Lee. Boosting the margin: A new explanation for the effectiveness of voting methods. *Annals of statistics*, 26(5): 1651–1686, 1998.

- Samarth Sinha, Homanga Bharadhwaj, Anirudh Goyal, Hugo Larochelle, Animesh Garg, and Florian Shkurti. Dibs: Diversity inducing information bottleneck in model ensembles. *arXiv:2003.04514*, 2020.
- Cedrique Tassi, Jakob Gawlikowski, Auliya Fitri, and Rudolph Triebel. The impact of averaging logits over probabilities on ensembles of neural networks. In *2022 Workshop on Artificial Intelligence Safety, AISafety 2022*, 2022.
- David MJ Tax, Robert PW Duin, and Martijn Van Breukelen. Comparison between product and mean classifier combination rules. In *Proc. Workshop on Statistical Pattern Recognition, Prague, Czech*, page 39. Citeseer, 1997.
- Robert Tibshirani. Bias, variance and prediction error for classification rules. Technical report, University of Toronto, 1996.
- Kagan Tumer and Joydeep Ghosh. Error correlation and error reduction in ensemble classifiers. *Connection science*, 8(3-4):385–404, 1996.
- Naonori Ueda and Ryohei Nakano. Generalization error of ensemble estimators. In *International Conference on Neural Networks (ICNN'96)*, volume 1. IEEE, 1996.
- Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *IEEE conference on computer vision and pattern recognition*, 2001.
- Andrew Webb, Charles Reynolds, Wenlin Chen, Henry Reeve, Dan Iliescu, Mikel Lujan, and Gavin Brown. To ensemble or not ensemble: When does end-to-end training fail? In *European Conference on Machine Learning*, pages 109–123, 2021.
- Danny Wood, Tingting Mu, and Gavin Brown. Bias-variance decompositions for margin losses. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- Yanzhao Wu, Ling Liu, Zhongwei Xie, Ka-Ho Chow, and Wenqi Wei. Boosting ensemble accuracy by revisiting ensemble diversity metrics. In *Conf. on Computer Vision & Pattern Recognition*, 2021.
- Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *Intl Conf. on Machine Learning*, 2020.
- Zhi-Hua Zhou. *Ensemble methods: foundations & algorithms*. Chapman & Hall/CRC, 2012.
- Zhi-Hua Zhou and Nan Li. Multi-information ensemble diversity. In *International Workshop on Multiple Classifier Systems*, pages 134–144. Springer, 2010.