

# Are Ensembles Getting Better All the Time?

**Pierre-Alexandre Mattei**

PIERRE-ALEXANDRE.MATTEI@INRIA.FR

*Université Côte d’Azur*

*Inria, Maasai team*

*Laboratoire J.A. Dieudonné, CNRS*

*Nice, France*

**Damien Garreau**

DAMIEN.GARREAU@UNI-WUERZBURG.DE

*Julius-Maximilians-Universität Würzburg*

*Institute for Computer Science / CAIDAS*

*Würzburg, Germany*

**Editor:** Maya Gupta

## Abstract

Ensemble methods combine the predictions of several base models. We study whether or not including more models always improves their average performance. This question depends on the kind of ensemble considered, as well as the predictive metric chosen. We focus on situations where all members of the ensemble are *a priori* expected to perform equally well, which is the case of several popular methods such as random forests or deep ensembles. In this setting, we show that ensembles are getting better all the time if, and only if, the considered loss function is convex. More precisely, in that case, the loss of the ensemble is a decreasing function of the number of models. When the loss function is nonconvex, we show a series of results that can be summarised as: ensembles of good models keep getting better, and ensembles of bad models keep getting worse. To this end, we prove a new result on the monotonicity of tail probabilities that may be of independent interest. We illustrate our results on a medical problem (diagnosing melanomas using neural nets) and a “wisdom of crowds” experiment (guessing the ratings of upcoming movies).

**Keywords:** Ensemble methods, deep ensembles, random forests, large deviations, wisdom of crowds

## 1. Introduction

The idea that groups are collectively wiser than individuals can be traced back to various branches of science, ranging from political science to philosophy, economics, or sociology. In statistics, forecasting, and machine learning, this idea is embodied by *ensemble methods*. Broadly speaking, ensembles combine the outputs of several base predictive models in a single decision, often by averaging (for regression) or by majority voting (for classification).

Early applications were mostly related to forecasting (Bates and Granger, 1969; Clemen, 1989) and decision theory (Stone, 1961; Genest and Zidek, 1986). In the 1990s, a significant amount of new ensembling techniques were developed by the burgeoning machine learning community, including notably boosting (Freund and Schapire, 1997), Breiman’s bagging (1996) or random forests (2001), Ho’s (1998) random subspace, neural network ensembles (Hansen and Salamon, 1990; Krogh and Vedelsby, 1994), and Bayesian model




Figure 1: Evolution of the accuracy of a dropout ensemble on the DermaMNIST dataset as the number of models grows (mean and standard deviation over 500 repetitions). *(Left)* The accuracy averaged over the whole test set has no clear monotonic pattern. *(Middle)* Accuracy averaged over images for which the asymptotic prediction (*i.e.* the prediction of an ensemble of infinite size)  $\bar{y}_\infty$  is correct. The accuracy is getting better when the ensemble size grows. *(Right)* Accuracy averaged over images for which  $\bar{y}_\infty$  is incorrect. The accuracy is going down. The behaviours of these middle and right panels are related to the non-convexity of the classification error, and are explained by our theory. Notice that the three  $y$ -axes have different scales.

averaging (Hoeting et al., 1999). Ensembles are routinely used in practice, for instance for probabilistic weather forecasting (Gneiting and Raftery, 2005) (see also Kuncheva’s (2014) monograph on the topic). Grinsztajn et al. (2022) go as far as claiming that random forests and boosting remain the most efficient machine learning methods for tabular data. More recently, ensembles of neural networks have been instrumental within the deep learning revolution, notably through dropout (Srivastava et al., 2014; Baldi and Sadowski, 2014; Hron et al., 2018) and deep ensembles (Lakshminarayanan et al., 2017) and their variants (see, *e.g.*, Sun et al., 2022; Laurent et al., 2023, and references therein). Random projections ensembles are still actively researched (Cannings and Samworth, 2017; Cannings, 2021).

The empirical successes of ensembles seem to indicate that the more models are being aggregated, the better the ensemble is. This, of course, can only be true on average. Adding a consistently bad model to the family will decrease the overall accuracy—for instance adding a model that always predicts the wrong class. The theoretical question we address in this paper is the following: **Is it always true that an ensemble of  $K + 1$  models performs better on average than an ensemble of  $K$  models?** We refer to this question as the *monotonicity question*. A positive answer would have practical consequences, encouraging us to always take more base models in our ensemble. As we will see, the picture is more nuanced and needs precise analysis to emerge.

### 1.1 A motivating example: neural nets for medical diagnosis

In this section, we consider the following practical problem: we want to predict whether a dermoscopic image corresponds to a benign keratosis-like lesion or a melanoma. We use




Figure 2: Evolution of the cross-entropy (or negative log-likelihood) of a dropout ensemble on the dermatology data set as the number of models grows (mean and standard deviation over 500 repetitions). *(Left)* Cross-entropy averaged over the whole test set. *(Middle)* Cross-entropy averaged over images for which the asymptotic prediction  $\bar{y}_\infty$  is correct. *(Right)* Cross-entropy averaged over images for which the asymptotic prediction  $\bar{y}_\infty$  is incorrect. As predicted by our theory, all curves are decreasing because the cross-entropy loss is convex. Notice that the three  $y$ -axes have different scales.

the DermaMNIST (Yang et al., 2023) data set, based on the HAM10000 collection (Tschandl et al., 2018), and retain only the classes “benign keratosis” and “melanoma” (more details on the data set are available in Section 5). Using a training set and a validation set, we train a LeNet-like convolutional neural network (LeCun et al., 1998) with dropout to discriminate these two classes. We then use a Monte Carlo dropout ensemble to classify the test set (Gal and Ghahramani, 2016; Hron et al., 2018), and look at the evolution of the accuracy (Figure 1) and the cross-entropy loss (Figure 2) as the ensemble grows. There is a crisp difference between the behaviours of the two losses. The overall accuracy has no clear monotonic pattern, but if we split the test set into two subsets depending on whether the asymptotic prediction is right or wrong, we notice that the accuracy is increasing for right predictions and decreasing for wrong ones! On the other hand, the cross-entropy appears to be clearly decreasing, regardless of the quality of the prediction. Our goal is to provide clear theoretical explanations for these phenomena.

## 1.2 Summary and organisation of the paper

In this work, we focus on the simplest instance of ensembles, that involve unweighted averaging of predictions that are “not too dependent.” While simple, this setting remains the usual formalization of most “wisdom of crowds” studies, and the basis of state-of-the-art machine learning methods like random forests and deep ensembles. In this context, formalised in Section 2, for an ensemble to be monotonically improving is deeply connected to the local geometry of the chosen loss function, more specifically to its convexity and its smoothness. More precisely, we show that:

- For **convex loss functions and exchangeable predictions** (Section 3), the expected loss is a non-increasing function of the number of ensemble members (Theo-

rem 2), generalising several results previously obtained for specific losses (Noh et al., 2017; Probst and Boulesteix, 2018). This result can be strengthened for strictly convex losses (Theorem 4).

- For **nonconvex loss functions and independent and identically distributed (i.i.d.) predictions** (Section 4), we show that, when the number of base models is large enough, good ensembles keep getting better, and bad ensembles keep getting worse in the following sense:
  - when the loss is sufficiently smooth, ensembles keep eventually getting better when the loss function is locally convex at their asymptotic predictions, whereas they keep getting worse if the loss is locally concave there (Theorem 5);
  - for the classification error, ensembles keep eventually getting better when their asymptotic prediction corresponds to the true label, and keep getting worse if it corresponds to a wrong label (Theorem 10). This result comes with not-so-mild regularity conditions that we relate to classical counterexamples *à la* Condorcet.
- To prove our result for nonsmooth losses, we prove two new theorems on the monotonicity of tail probabilities that may be of independent interest (Theorems 7 and 8). This result is based on “strong” large deviations theorems (Bahadur and Ranga Rao, 1960; Petrov, 1965; Joutard, 2017).

Informally, our results may be summarised by the sentence **ensembles are getting better all the time if, and only if, the considered loss function is convex**. We support these claims by two real data experiments presented in Section 5.

### 1.3 Related works: monotonicity and the wisdom of crowds

The main heuristic justification of ensembles is that groups are collectively wiser than individuals, because aggregated decisions have a smaller variance than individual ones. This general idea is referred to in various, sometimes colorful, ways: some, like Dekel and Shamir (2009), praise the “vox populi,” following Galton’s (1907) seminal paper; the phrase “wisdom of crowds,” used for instance by Prelec et al. (2017) or Simoiu et al. (2019), was immensely popularised by Surowiecki’s (2004) bestselling book; political scientists like Althaus (2003) often debate the virtues of the “miracle of aggregation” as a theoretical argument in favor of democracy; many scientists made the popular phrase “two heads are better than one” their own (*e.g.*, Koriati, 2012; Chu-Carroll et al., 2003). A recent review of the philosophical ramifications of collective intelligence was provided by Elkin and Pettigrew (2025).

Condorcet (1785) proposed a first mathematical formalisation of this heuristic insight. His goal was to prove that, within a trial, the larger the number of jurors, the better. In modern machine learning jargon, he showed that ensembles containing an odd number of Bernoulli-distributed base models are monotonically improving, with respect to the classification error (see Shteingart et al. 2020, for a modern proof). Surprisingly, he also noticed that monotonicity is broken if the ensemble is allowed to contain an even number of models.

An important (and quite diverse) amount of work has been devoted to the theoretical study of ensembles; while giving a detailed overview is out of our scope, it is worth providing

a few entry points to this literature. Some authors analysed the properties of specific kinds of ensembles, such as bagging (Larsen, 2023), random forests (Scornet and Hooker, 2025), random projections (Cannings and Samworth, 2017), dropout (Hron et al., 2018) or deep ensembles (Wild et al., 2023). Others studied the relationship between the benefits of ensembling and the diversity of the ensemble (Ortega et al., 2022; Theisen et al., 2023). More related to our work are papers that study the behaviour of ensembles when  $K$  is large. In particular, Cannings and Samworth (2017) and Lopes (2020) derived asymptotic expansions of the test loss when  $K \rightarrow \infty$ . Unfortunately, such expansions do not capture monotonicity, even for large enough  $K$  (this will be discussed further in Section 4.2).

However, in spite of this large body of work, the monotonicity question has been quite overlooked, especially theoretically. Lam and Suen (1997), Berend and Sapir (2005), and Dougherty and Edward (2009) did revisit Condorcet’s theorem, and studied the issues that make the presence of an even number of models problematic. Several generalisations of this theorem were also derived, in particular with weaker dependence assumptions (see, *e.g.*, Ladha, 1993). Closer to our investigations Probst and Boulesteix (2018), in the context of random forests, performed a very thorough theoretical and empirical study of the monotonicity question. Their empirical conclusions were that, for some performance metrics (the cross-entropy or the Brier score for classification and the mean squared error for regression), ensembles are getting better all the time, but that is not necessarily the case for the classification error. Moreover, they provided theoretical support for some of these phenomena. Their paper was a key inspiration for this work, and our main contribution is to provide a unified view of their insights by essentially showing that ensembles are getting better all the time if, only if, the considered loss function is convex. Indeed, the cross-entropy, the Brier score, and the mean squared error are all convex losses (and always lead to monotonic improvements), whereas the classification error is not (and may lead to non-monotonicity).

The relationship between the usefulness of ensembling and convexity was perhaps first highlighted by Stone (1961) in the context of decision theory. It was further discussed by several papers in the 1990s (McNees, 1992; Perrone, 1993; Breiman, 1996; George, 1999). An interesting historical perspective on the question was provided by Manski (2010, Section 6). However, the relationship between the monotonicity question and convexity appears to have been overlooked.

**Monotonicity beyond ensembles** It is worth mentioning some related work on monotonicity in statistics and machine learning. A first intriguing phenomenon is the fact that the coverage probability of a confidence interval for the true parameter of a Bernoulli random variable is non-monotonic (see, *e.g.*, Brown et al. 2001, Figure 1). In plain words, and perhaps counter-intuitively, increasing the number of observations does not guarantee that the probability for the confidence interval to contain the true parameter increases! Another related question is the one of the *monotonicity of learning*: will a machine learning algorithm perform better when adding one new sample to the training set? Devroye et al. (2013) call such class of models *smart rules* predictors and notices that it is quite challenging to know for sure which models are “smart.” Indeed, most consistency results only provide an asymptotic statement in the sample size, with no monotonicity guarantee. More recently, this problem was cast as an open problem by Viering et al. (2019) at the 32nd Conference on Learning Theory (COLT). Progress on that front was reviewed by Viering and Loog

(2023). Finally, the question of monotonicity also arises when looking at the variance of a model, as illustrated, for instance by Rasmussen and Williams (2006, Exercise 2.4), for Gaussian processes.

## 2. Background, notations, and examples

### 2.1 A general predictive framework

We consider a predictive task where several methods output predictions  $\hat{y}$ , living in a convex set  $C$ . The quality of these predictions is evaluated using a loss function  $L : C \rightarrow \mathbb{R}$ . This broad setting contains several important cases, and we present below a few important examples. In most of these examples,  $\hat{y}$  will be of the form  $f(x)$ , where  $f$  is a base model and  $x$  a given fixed data point. This means that  $L(\hat{y})$  is the loss of an *individual prediction*, rather than the loss averaged over a dataset. Note that we require the prediction set to be convex because we want that averaged predictions remain inside  $C$  (this is barely a restriction, since when the prediction set is nonconvex we can simply replace it by its convex hull).

**Supervised learning and forecasting** are the main examples we are interested in. Consider a given data point  $x \in \mathcal{X}$ , with true label/response  $y \in \mathcal{Y}$ . The prediction will typically be  $\hat{y} = f(x)$  where  $f$  can be for instance a neural network, a random forest, or even a human expert opinion (*e.g.*, a doctor diagnosing a patient based on their records  $x$ ). In this setting, loss functions have the familiar form  $L(\hat{y}) = \ell(\hat{y}, y)$  where  $\ell : C \times \mathcal{Y} \rightarrow \mathbb{R}$  measures the mismatch between the true label and its prediction (see, *e.g.*, Shalev-Shwartz and Ben-David 2014). As an example, for neural net classification with  $n_{\text{cl}}$  classes, the prediction set  $C$  will either be  $\mathbb{R}^{n_{\text{cl}}}$  (in the absence of a final softmax layer, in this case  $\hat{y}$  denotes the logits) or the  $n_{\text{cl}}$ -simplex ( $\hat{y}$  will then be the output of the softmax layer). In both cases,  $\ell$  will typically be the cross-entropy, defined in the case where  $\hat{y}$  is the softmax output, as

$$L(\hat{y}) = \ell(\hat{y}, y) = \sum_{j=1}^{n_{\text{cl}}} -y_j \log \hat{y}_j, \quad (1)$$

where the true label is represented by a one-hot encoding  $y = (y_1, \dots, y_{n_{\text{cl}}}) \in \{0, 1\}^{n_{\text{cl}}}$ . Another common choice of loss function in that case is the classification error, also known as the 0 – 1 loss. Denoting  $\text{argmax}(u)$  the index of the coordinate of  $u$  with the largest value, the classification error can be written

$$L(\hat{y}) = \ell(\hat{y}, y) = \begin{cases} 1 & \text{if } \text{argmax}(\hat{y}) \neq \text{argmax}(y) \\ 0 & \text{if } \text{argmax}(\hat{y}) = \text{argmax}(y). \end{cases} \quad (2)$$

If the response set is continuous (*e.g.*, if  $\mathcal{Y} = \mathbb{R}$  as in a standard regression problem), the prediction set can be  $\mathbb{R}$  if we are simply interested in predicting a point estimate of the response, then  $\ell$  can be the squared error:  $L(\hat{y}) = \ell(\hat{y}, y) = (\hat{y} - y)^2$ . If we are interested in the full conditional distribution of the response,  $C$  can be defined as the set of probability densities over  $\mathbb{R}$ . The loss  $\ell$  can then be any scoring rule like the negative log-likelihood  $-\log \hat{y}(y|x)$  or the continuous ranked probability score, see, *e.g.*, Gneiting and Raftery (2007). Other supervised examples with more exotic prediction spaces and loss functions include image segmentation or ordinal regression.

**Density estimation** is generally unsupervised, and not directly a predictive problem. It is nevertheless a particular case of our framework. Here,  $\hat{y}$  will be a probability density over the data space  $\mathcal{X}$ , for instance modelled by a mixture model or a deep generative model, which means that  $C$  will be the set of all densities over  $\mathcal{X}$ . For a given  $x \in \mathcal{X}$ , a popular loss function is then the negative log-likelihood  $L(\hat{y}) = -\log \hat{y}(x)$ . Similarly to the probabilistic regression example outlined above, any scoring rule can also be chosen (*e.g.*, score matching and its variants). More exotic losses include using statistical divergences between the distribution  $\hat{y}$  and the empirical distribution of a test or validation set. While ensembling is less common for density estimation than forecasting and supervised learning, it has been successfully used in several contexts. For instance, Riesselman et al. (2018) used ensembles of deep generative models to model biological sequences and Russell et al. (2015) and Casa et al. (2021) have used ensembles of Gaussian mixtures for density estimation and clustering.

**Statistical estimation** involves an unknown parameter  $\theta^* \in \Theta$ . In this case,  $C = \Theta$  and  $\hat{y}$  can be any estimator of  $\theta$ . A possible loss function in this example is given by  $L(\hat{y}) = d(\hat{y}, \theta^*)$ , where  $d$  is a distance over  $\Theta$ . Here again, the use of ensembles is less common than for supervised learning. Nonetheless, bagging and its variants have been thoroughly studied for statistical estimation (see, *e.g.*, Bühlmann and Yu 2002; Friedman and Hall 2007; Bach 2008).

**Why is  $L$  only a function of  $\hat{y}$ ?** In many of these examples, the loss function depends on other quantities than the predictions (the features, the true label...), yet we just denote it by  $L(\hat{y})$ . Our analysis aims at studying the impact of the predictions on the loss, therefore we keep all these other quantities fixed, and vary only the predictions. This is why we only model our loss function  $L$  as a function of the predictions  $\hat{y}$ . In particular, this means that our results typically concern the loss evaluated for a single given data point and a single ground truth label. Of course, we are also often interested in evaluating the loss not only on individual data points, but on batches or entire data sets. Our results on individual data points can then be extended by averaging them, as we discuss further in Section 2.2.3.

## 2.2 Several predictions, and how to average them

From now on, we consider that we are given  $K \in \mathbb{N}^*$  different predictions, named  $\hat{y}_1, \dots, \hat{y}_K \in C$ , that we wish to aggregate.

### 2.2.1 WHERE DO OUR PREDICTIONS $\hat{y}_1, \dots, \hat{y}_K$ COME FROM?

A fundamental example to keep in mind is the situation where the predictions come from slight variations from the same model. In Breiman’s (1996) *bagging*, each  $\hat{y}_k$  comes from using the same algorithm on a different bootstrapped version of the original data. This is also the cornerstone of random forests, for which each  $\hat{y}_k$  is the output of a single (random) decision tree (Breiman, 2001). Each tree has a different structure coming from the randomness of the tree construction procedure, leading to distinct, random  $\hat{y}_k$ s. Another important example, pioneered by Hansen and Salamon (1990) and popularised by Lakshminarayanan et al. (2017), uses as  $\hat{y}_k$ s the outputs of neural networks trained with different weights initialisations. Another popular way of creating ensembles of deep nets is via *dropout* (Srivastava et al., 2014; Baldi and Sadowski, 2014): each  $\hat{y}_k$  is obtained by sampling a different dropout

mask and applying it to the same backbone network. Gal and Ghahramani (2016) call such a technique Monte Carlo dropout.

In all the examples presented in the previous paragraph, the  $\hat{y}_1, \dots, \hat{y}_K$  are i.i.d. random variables. Assumptions of this sort will be required for our proofs. Indeed, with the exception of Lemma 1, we will always require our predictions to be random variables with some (weak) dependence structure. We refer to Section 3.3 for a simple example of what can go wrong when there is too much dependency between the individual models.

We want to emphasise that our framework also encompasses situations where we are given a few fully deterministic predictions (for instance, the forecasts of several different physical models). One way to see them as random predictions and to make them amenable to our theory is to assume that the ordering of the  $\hat{y}_1, \dots, \hat{y}_K$  does not matter, or equivalently that we have no knowledge of how these models fare against each other. This can be formalised by assuming that we do not observe  $\hat{y}_1, \dots, \hat{y}_K$  but  $\hat{y}_{\sigma(1)}, \dots, \hat{y}_{\sigma(K)}$ , where  $\sigma$  is a uniformly sampled permutation of  $\{1, \dots, K\}$ . Because  $\sigma$  is random,  $\hat{y}_{\sigma(1)}, \dots, \hat{y}_{\sigma(K)}$  is also random. While not i.i.d. in general, this *randomly reordered ensemble* satisfies another weaker condition called *exchangeability*. A sequence of random variables is exchangeable when any of its reorderings has the same distribution, we will give a formal definition in Eq. (13).

### 2.2.2 HOW TO AGGREGATE PREDICTIONS?

Perhaps the simplest idea to aggregate predictions is to take their (unweighted) mean, namely to consider the prediction

$$\bar{y}_K = \frac{1}{K} \sum_{k=1}^K \hat{y}_k. \quad (3)$$

Note that we are allowed to consider such an aggregation since  $C$  is convex, which implies that the ensemble prediction  $\bar{y}_K$  indeed belongs to  $C$ . This simple aggregation scheme is used by many ensembles, for instance bagging (and in particular random forests), deep ensembles, or Monte Carlo dropout. Many other aggregation schemes have been proposed (see, *e.g.*, Kucheva 2014, Chapters 4 and 5), but our analysis is fundamentally tied to the simple averaging technique, so we will not discuss them further. A convenient property of  $\bar{y}_K$  is that its asymptotic behaviour is quite intuitive. Indeed, when  $\hat{y}_1, \dots, \hat{y}_K$  are i.i.d. with mean  $\bar{y}_\infty = \mathbb{E}[\hat{y}_1] = \dots = \mathbb{E}[\hat{y}_K]$ , then the law of large numbers implies that  $\bar{y}_K \rightarrow \bar{y}_\infty$  when  $K \rightarrow +\infty$ . We will refer to  $\bar{y}_\infty$  as the *asymptotic prediction* of the ensemble.

Several other aggregation schemes can, in some settings, be equivalent to the simple average. For instance, the geometric mean can be written as a regular mean in log scale, since  $\sqrt[K]{\hat{y}_1 \cdots \hat{y}_K} = \exp\left(\frac{1}{K} \sum_{k=1}^K \log \hat{y}_k\right)$ . The popular method of *majority voting* can also be seen as a form of standard average. Indeed, consider for instance a binary classification problem for which several votes  $\hat{y}_1, \dots, \hat{y}_K \in \{0, 1\}$  are cast for a true class  $y \in \{0, 1\}$ . The prediction space  $\{0, 1\}$  is nonconvex, but we may choose to see the votes as elements of its convex hull  $C = [0, 1]$ . A natural choice of loss function is then the classification error  $L(\hat{y})$ , defined by Eq. (2). The most natural way of combining these votes is to use majority voting, considering as final choice  $\tilde{y}_{\text{MV}, K} = \mathbf{1}(\frac{1}{K} \sum_{k=1}^K \hat{y}_k > 0.5)$ . It turns out this majority



vote will have the same loss as the ensemble prediction:  $L(\tilde{y}_{\text{MV},K}) = L(\bar{y}_K)$ , which means that majority voting is equivalent to simple averaging.

### 2.2.3 WHAT SORT OF RESULTS ARE WE GOING TO PROVE?

A natural way of looking at the performance of the random ensemble prediction is to study its expectation. Our goal will be to prove that, under various settings,  $\mathbb{E}[L(\bar{y}_K)]$  is monotonic, where  $\mathbb{E}$  is taken with respect to the randomness associated to the predictions. Sometimes, we will merely show that  $\mathbb{E}[L(\bar{y}_K)]$  is *eventually monotonic*, which means that there exists an integer  $K_0$  such that  $(\mathbb{E}[L(\bar{y}_K)])_{K \geq K_0}$  is monotonic.

A few comments about what this means are in order. The quantity  $\mathbb{E}[L(\bar{y}_K)]$  corresponds to the loss of a *single* prediction problem, *e.g.*, predicting the label of a single image, averaged over the randomness of the ensemble. This is in contrast with many machine learning results that average over a dataset. Nevertheless, the monotonicity of individual losses will imply monotonicity for losses averaged over datasets by simply applying our results to every single point of the dataset and using the fact that an average of decreasing (respectively increasing) sequences remains decreasing (respectively increasing). This is illustrated in Fig. 2: our results imply that the cross-entropy loss decreases for every single test point, and in turn that the whole test cross-entropy decreases. Things are slightly different in Fig. 1: we cannot average our results over the whole test set (since it is composed of points whose accuracy eventually grows and points whose accuracy eventually decreases). However, we can split the test dataset into two sub-datasets for which we can apply our theorems, leading to the middle and right panels of Fig. 1.

## 3. Ensembles get better for convex losses

In this section, the loss function  $L : C \rightarrow \mathbb{R}$  is always assumed to be convex. In a supervised context where  $L(\hat{y}) = \ell(\hat{y}, y)$ , this means that  $\ell$  is convex in its first argument. In classification, this is for instance the case of the cross-entropy, the Brier score, or the hinge loss. In regression, convex losses include the mean absolute error, the mean squared error, or Huber’s (1964) robust loss. Examples in image segmentation include the cross-entropy or the Lovász-softmax loss (Berman et al., 2018). Another ubiquitous example of convex loss is the negative log-likelihood, both for density estimation and probabilistic prediction. In decision theory, convex losses correspond to concave utilities and to risk aversion (see *e.g.* Parmigiani and Inoue 2009, Chapter 4).

### 3.1 Warm-up: An ensemble predicts better than a single model

Let us begin by reviewing a few classical arguments in favour of ensembling with a convex loss, based on Jensen’s inequality. The simple and compelling reasoning below, that dates back to early works on forecasting at least (Winkler, 1971, Section 3.4), was called the “algebraic wisdom of crowds” by Manski (2016). Similar arguments were also given, for instance, by McNeis (1992) or Breiman (1996, Section 4) in the specific cases of the squared and absolute losses and by Perrone (1993) and George (1999) for a generic convex loss. The

convexity of  $L$  directly implies that, for all  $\hat{y}_1, \dots, \hat{y}_K \in C$ ,

$$L(\bar{y}_K) \leq \frac{1}{K} \sum_{k=1}^K L(\hat{y}_k) \leq \max_{k \in \{1, \dots, K\}} L(\hat{y}_k). \quad (4)$$

Therefore the ensemble will perform better, on average, than any of its components randomly chosen (with uniform probability  $1/K$ ). In particular, the ensemble will *always* perform better than its worst component.

When all elements of the ensemble are identically distributed, so are  $L(\hat{y}_1), \dots, L(\hat{y}_K)$  and integrating the left-hand side of Eq. (4) leads to

$$\mathbb{E}[L(\bar{y}_K)] \leq \mathbb{E}[L(\hat{y}_1)] = \dots = \mathbb{E}[L(\hat{y}_K)], \quad (5)$$

assuming that all involved expectations exist. This means that an ensemble of identically distributed predictors performs better than a single one of them, on average. If we further assume that the ensemble has mean  $\bar{y}_\infty = \mathbb{E}[\hat{y}_1] = \dots = \mathbb{E}[\hat{y}_K]$ , then

$$L(\bar{y}_\infty) \leq \mathbb{E}[L(\bar{y}_K)], \quad (6)$$

due to Jensen's inequality and the fact that  $\mathbb{E}[\hat{y}_k] = \bar{y}_\infty$  for all  $k$ . We chose to denote this mean  $\bar{y}_\infty$  because the law of large numbers implies that  $\bar{y}_K \rightarrow \bar{y}_\infty$  almost surely when  $K \rightarrow \infty$ , if we further assume that the  $\hat{y}_k$ s are independent. According to Eq. (6), this “infinite” ensemble that would be obtained by averaging over all possible predictions, is better than any of its finite counterparts. However, this does not say how these finite ensembles fare against each other. As we shall see now, this question can also be addressed by playing around with Jensen's inequality.

### 3.2 A first monotonicity lemma

All the new results in this section are based on a simple identity, that expresses  $\bar{y}_K$  as a convex combination of smaller ensembles. The starting point is to notice the following equality between  $K$ -dimensional vectors:

$$\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \\ 1 \end{bmatrix} = \frac{1}{K-1} \left( \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 1 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 1 \\ 1 \end{bmatrix} + \dots + \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 0 \\ 1 \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \\ 0 \end{bmatrix} \right). \quad (7)$$

This directly leads to

$$\frac{1}{K} \sum_{k=1}^K \hat{y}_k = \frac{1}{K} \sum_{j=1}^K \frac{1}{K-1} \sum_{k \neq j} \hat{y}_k, \quad (8)$$

which, combined with the convexity of  $L$ , leads to the following lemma.

**Lemma 1 (monotonicity lemma).** *Assume that  $L$  is a convex function. Then, for any  $\hat{y}_1, \dots, \hat{y}_K \in C$ , we have*

$$L\left(\frac{1}{K} \sum_{k=1}^K \hat{y}_k\right) \leq \frac{1}{K} \sum_{j=1}^K L\left(\frac{1}{K-1} \sum_{k \neq j} \hat{y}_k\right). \quad (9)$$

This inequality may already be interpreted as a monotonicity result: removing an element of the ensemble uniformly at random will hurt, on average. Note that this comes *without any assumption on the ensemble itself*, which could be composed of very different and potentially deterministic models. To prove stronger results, we need to make additional assumptions on the ensemble. Our assumptions will be probabilistic, either requiring independent, identically-distributed predictions, or a weakened version of this assumption called *exchangeability*. They will translate the fact that we suspect that, *a priori*, all elements of the ensembles perform equally well.

### 3.3 Exchangeable ensembles are getting better all the time

We saw in the previous section (Eq. 6) that assuming that the ensemble is identically distributed led to an asymptotic form of monotonicity: the infinite ensemble  $\bar{y}_\infty$  performs better than any finite ensemble, which in turn performs better than any single model, because of Eq. (5). However, the identically distributed assumption alone will not take us very far. Indeed, let us consider the following example: an ensemble of three models such that  $\hat{y}_1$  and  $\hat{y}_2$  are independent and identically distributed, and  $\hat{y}_3 = \hat{y}_1$ . Then, one can show that the ensemble gets *worse* in average when adding its third component:

$$\mathbb{E} \left[ L \left( \frac{\hat{y}_1 + \hat{y}_2}{2} \right) \right] \leq \mathbb{E} \left[ L \left( \frac{\hat{y}_1 + \hat{y}_2 + \hat{y}_3}{3} \right) \right]. \quad (10)$$

This directly follows from the identity

$$\frac{\hat{y}_1 + \hat{y}_2}{2} = \frac{1}{2} \left( \frac{2\hat{y}_1 + \hat{y}_2}{3} + \frac{\hat{y}_1 + 2\hat{y}_2}{3} \right), \quad (11)$$

combined with the Jensen's inequality and the fact that

$$\hat{y}_1 + \hat{y}_2 + \hat{y}_3 = 2\hat{y}_1 + \hat{y}_2 =_d 2\hat{y}_2 + \hat{y}_1, \quad (12)$$

where  $=_d$  denotes equality in distribution.

This simple counter-example motivates the use of a stronger assumption than identically distributed. A natural choice would be to assume that the elements of the ensembles are i.i.d., which is often the case in practice (for instance in random forests or deep ensembles). We will rely on an assumption that is actually weaker than i.i.d. but stronger than i.d.: *exchangeability*. We remind that a sequence of random variables  $\hat{y}_1, \dots, \hat{y}_K$  is exchangeable when, for any permutation  $\sigma \in S_K$  of the indices,

$$(\hat{y}_1, \dots, \hat{y}_K) =_d (\hat{y}_{\sigma(1)}, \dots, \hat{y}_{\sigma(K)}). \quad (13)$$

For more insights on exchangeability, we recommend Diaconis (1988). Intuitively, an ensemble is exchangeable when the ordering of its members is not important.

Many popular ensembles are exchangeable. This is of course the case of all i.i.d. ensembles (including bagging, random forests, random projections, deep ensembles), but there also exists non-i.i.d. exchangeable ensembles. One example that was sketched at the end of Section 2.2.1 is the one of a randomly reordered ensemble. Another general example is the one of cross-validated committees (Rokach, 2010, Section 2.2.4) that train different

models on different exchangeables subsets of the training data. Political scientists have also argued that exchangeability is a more sensible assumption than i.i.d. for jury theorems *à la* Condorcet (Ladha, 1993; Peleg and Zamir, 2012).

However, there are also ensembles for which the exchangeability assumption is inherently wrong. This is notably the case of boosting, or any similar technique that trains models iteratively, making the construction of the  $k$ th base model dependent on the  $k - 1$  previous models, and hereby making their ordering essential. This is also the case when we have prior knowledge that some models are much better than others. In these two cases, using a weighted average instead of just the usual mean  $\bar{y}_K$  is more advisable.

Under this assumption of exchangeability, we can show the following result, that has a simple interpretation: *when the loss is convex and the ordering does not matter, ensembles monotonically get better, on average.*

**Theorem 2 (monotonicity for convex losses, Marshall and Proschan, 1965).** *If the loss  $L$  is convex and  $\hat{y}_1, \dots, \hat{y}_K$  are exchangeable, then,*

$$\mathbb{E}[L(\bar{y}_K)] \leq \mathbb{E}[L(\bar{y}_{K-1})]. \quad (14)$$

**Proof** From Lemma 1, we have, by integrating

$$\mathbb{E}\left[L\left(\frac{1}{K}\sum_{k=1}^K\hat{y}_k\right)\right] \leq \frac{1}{K}\sum_{j=1}^K\mathbb{E}\left[L\left(\frac{1}{K-1}\sum_{k\neq j}\hat{y}_k\right)\right]. \quad (15)$$

But, because  $\hat{y}_1, \dots, \hat{y}_K$  are exchangeable, for any  $j \in \{1, \dots, K\}$ ,

$$\mathbb{E}\left[L\left(\frac{1}{K-1}\sum_{k\neq j}\hat{y}_k\right)\right] = \mathbb{E}\left[L\left(\frac{1}{K-1}\sum_{k=1}^{K-1}\hat{y}_k\right)\right], \quad (16)$$

thus all the terms of the sum of the right-hand side of Eq. (15) are equal to the right-hand side of Eq. (14), which concludes the proof.  $\blacksquare$

A version of Theorem 2 was first proved by Marshall and Proschan (1965) in a very different context, using the concept of *majorisation* (see, *e.g.*, Marshall et al. 2011). In the machine learning community, a similar result was independently rediscovered and popularised in the context of variational inference for deep generative models by Burda et al. (2016). The connection between majorisation and the results of Marshall and Proschan (1965) and Burda et al. (2016) was highlighted and discussed by Mattei and Frellsen (2022). Struski et al. (2023) provided an alternative proof of the result of Burda et al. (2016) that is similar to our short proof of Theorem 2.

None of the works mentioned in the previous paragraph are related to ensembles. However, in the context of ensembles, several particular cases of Theorem 2 for specific loss functions have been previously proved. In particular, Probst and Boulesteix (2018) dealt with two cases: when  $L$  is the square loss and when  $L$  is the cross entropy and  $K$  is large enough. Moreover, Noh et al. (2017) noticed that the result of Burda et al. (2016) could be applied to ensembles when  $L$  is the negative log-likelihood.

### 3.4 Strict improvements for strongly convex losses

The previous result show that the average loss is non-increasing. Under which conditions is it *strictly* decreasing? This depends both on the properties of the loss (for instance, if  $L$  is constant, no ensemble can strictly improve) and of the ensemble (an ensemble containing identical copies of the same predictor will have constant loss). We will see now that two simple assumptions are sufficient to ensure strict improvements: strong convexity of the loss, and existence of a nonsingular covariance of the ensemble. In this subsection, we further assume that the prediction set  $C$  is equipped with a norm  $\|\cdot\|$ .

**Definition 3 (strong convexity).** *Let  $\mu > 0$ . A function  $L : C \rightarrow \mathbb{R}$  is  $\mu$ -strongly convex when, for all  $x, y \in C$ , we have*

$$L(tx + (1-t)y) \leq tL(x) + (1-t)L(y) - \mu t(1-t)\|x - y\|^2. \quad (17)$$

Strong convexity leads to a strengthening of Jensen’s inequality (see, *e.g.*, Nikodem 2014, Theorem 2) that we can combine with our key identity Eq. (8) to prove strict improvements for i.i.d. ensembles.

**Theorem 4 (strict monotonicity for strongly convex losses).** *Assume that the ensemble is exchangeable and has second order moments. Let us write  $\Sigma \succ 0$  for its covariance matrix. Define  $\rho := \text{tr}(\text{Cov}(\hat{y}_1, \hat{y}_2))/\text{tr}(\Sigma)$  its correlation coefficient. Then, if the loss  $L$  is  $\mu$ -strongly convex, then*

$$\mathbb{E}[L(\bar{y}_K)] \leq \mathbb{E}[L(\bar{y}_{K-1})] - \frac{\mu}{K(K-1)}(1-\rho)\text{tr}(\Sigma). \quad (18)$$

The proof of Theorem 4 can be found in Appendix A. For i.i.d. ensembles like random forests of deep ensembles, we have  $\rho = 0$  which slightly simplifies the bound. This bound is tight when  $L$  is the square loss and  $\hat{y}_1, \dots, \hat{y}_K$  are i.i.d. standard Gaussian. When  $\hat{y}_1 = \dots = \hat{y}_K$ ,  $\rho = 1$ , and the bound is also tight.

This result formalises the fact that we can expect the loss improvements of ensembling to be larger as the variance of the ensemble grows. One may wonder what happens in the extreme case where the variance is infinite. Perhaps surprisingly, nothing clear can be said in this case. Indeed, take for instance  $\hat{y}_1, \dots, \hat{y}_K$  to be i.i.d. Cauchy random variables. Then,  $\bar{y}_K$  is also Cauchy distributed, which means that the performance of the ensemble will be constant for any loss function, even though the variance is infinite.

## 4. What happens for nonconvex losses?

For convex losses, we saw that the theory is quite simple and compelling: the average performance of ensembles is always improving. The picture becomes much more complex when the loss is no longer assumed to be convex. Indeed, we saw in Section 1.1 that there was a crisp difference between the cross-entropy loss and the accuracy.

As we will see, there is nonetheless a general pattern for nonconvex losses. Some ensembles get better: they are the ones that, on average, “get it right” (in the sense that  $\bar{y}_\infty$  is a “good prediction”), and correspond for instance to the ones on the middle panel of Fig. 1. On the other hand, when  $\bar{y}_\infty$  is a “bad prediction,” we will see that ensembling hurts (as it does on the right panel of Fig. 1).

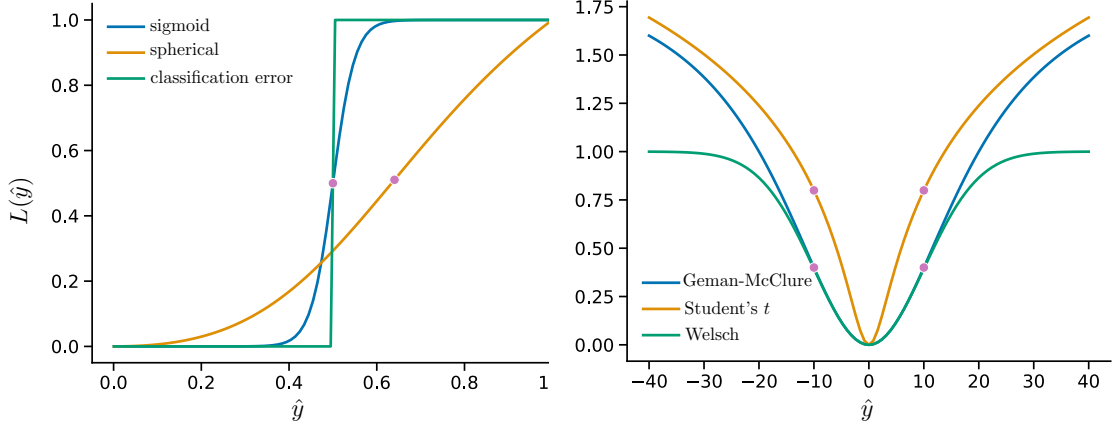


Figure 3: Nonconvex loss functions for binary classification (left panel) and regression (right panel). In both case the true label/response is  $y = 0$ . For smooth losses, purple dots denote inflexion points: the frontier between the concave part of the loss (that corresponds to “right” predictions) and the concave one (that corresponds to “wrong” ones).

There are two sorts of nonconvex loss that will interest us. The first ones are the sufficiently smooth ones. Their monotonicity will be analysed by Taylor-expanding the loss around  $\bar{y}_\infty$ . We will then study the most popular nonconvex loss: the classification error, that will prove challenging to tackle due to its discontinuity. Unlike results for the convex case, which were true for any  $K$ , all results in this section will be asymptotic, in the sense that we will prove that  $\mathbb{E}[L(\bar{y}_K)]$  is *eventually monotonic* (i.e., that it is monotonic for  $K$  large enough).

#### 4.1 Smooth nonconvex losses

A first case of nonconvex loss that is easy to tackle is the one of a *concave* loss. Indeed, if  $L$  is concave, then  $-L$  is convex which, by virtue of Theorem 2, implies that  $\mathbb{E}[-L(\bar{y}_K)]$  decreases and thus that  $\mathbb{E}[L(\bar{y}_K)]$  increases: for a concave loss, ensembles are getting worse all the time! However, to the best of our knowledge, concave losses are never used in statistics and machine learning. Laan et al. (2017) considered the example of a “square-root cost”  $L(\hat{y}) = \sqrt{|y - \hat{y}|}$  which is locally concave everywhere except at zero, but not *globally* concave (which prevents us from using Theorem 2).

Although globally concave loss are not used, *some losses are locally concave in some part of the prediction space*, and locally convex on another. When these losses are smooth enough, these concave/convex parts can be recognised by checking if the Hessian of the loss is negative/positive definite.

Most of these smooth nonconvex loss functions share the informal property that *the loss is locally convex in parts of the prediction space where predictions are “right,” and locally concave in parts where they are “wrong”*. We have remained so far voluntarily vague

regarding what we meant by “right” and “wrong,” because these terms will have different precise meanings depending on the loss functions. To illustrate this general behaviour, let us look at a few examples.

#### 4.1.1 EXAMPLES OF SMOOTH NONCONVEX LOSSES

We begin with losses designed for classification and consider for simplicity the binary case. We assume that the correct class is  $y = 0$ , and the prediction can be either an estimated probability  $\hat{y} \in [0, 1]$  that 1 is the correct class, or a real-valued score  $\hat{y} \in \mathbb{R}$  that is higher when the classifier is more confident that the label is 1. We describe below three recipes for designing nonconvex smooth losses, and plot examples in Fig. 3.

- **Smooth approximations of the classification error** have been used as surrogates of it, and are generally nonconvex. An important instance is the *sigmoid loss*. This sigmoid loss was popularised by its use within the DOOM II algorithm of Mason et al. (1999). It was also notably studied by Bartlett et al. (2006, Example 7), who showed it was classification-calibrated. This loss is easier to interpret in the probabilistic prediction setting, when  $\hat{y} \in [0, 1]$ . Indeed, in this case, the sigmoid loss is locally convex when  $\hat{y} < 0.5$ , *i.e.*, when the prediction using a 50% threshold is correct, and locally concave when  $\hat{y} > 0.5$ , *i.e.* when the prediction using the same threshold is wrong. The *Savage loss* introduced in a boosting context by Masnadi-Shirazi and Vasconcelos (2008) and the *normalised cross entropy* used by Ma et al. (2020) to train deep nets with noisy labels have a similar sigmoidal shape. They also share the same property of being convex when  $\hat{y} < 0.5$ , and concave when  $\hat{y} > 0.5$ . For these losses, a “right prediction” is one whose most probable class is the correct one. In an opposite fashion, a “wrong prediction” gives the highest probability to the wrong class.
- **Scoring rules** are the most common way to design losses in the probabilistic setting (when  $\hat{y} \in [0, 1]$ ). Most strictly proper scoring rules for classification are convex (*e.g.* the cross-entropy or the Brier score), with the notable exception of the *spherical score* (see, *e.g.*, Gneiting and Raftery, 2007). Like the sigmoid loss, the spherical score is convex when  $\hat{y}$  is small and concave when  $\hat{y}$  is large. However, unlike the sigmoid whose inflection point was at  $\hat{y} = 50\%$ , the inflection point of the spherical score is at  $\hat{y} = 1/8 + \sqrt{17}/8 \approx 64\%$ . We give a formal definition in Appendix B. For the spherical loss, a “right prediction” is thus one that will predict the true label  $y = 0$  using the somewhat peculiar cutoff of 64% instead of the usual 50%. This illustrates the asymmetry of the spherical score.
- One way to avoid overfitting is to use a loss function that makes the model pay a price for overconfident predictions (potentially *even correct ones*). This is the rationale behind the **tangent loss**, introduced by Masnadi-Shirazi et al. (2010) as a boosting objective. Since this is not a probabilistic loss, predictions are real-valued scores  $\hat{y} \in \mathbb{R}$ . The tangent loss is concave when  $\hat{y}$  is either too large or too small, penalising both overconfident mistakes and successes. In this case, “right predictions” correspond to the ones that are not overconfident. For more details on the tangent loss, see Appendix B.

Regarding regression, the most commonly occurring nonconvex smooth losses are used to induce robustness to outliers. The main motivation for such losses is that the squared loss will typically make the model pay an extremely large price for large errors. A way to fix this behavior is to design losses that resemble the squared error when errors are small, but grow less steadily when errors are large. The most famous robust loss is Huber’s (1964), which is equal to the squared error when the error is below a threshold, and grows linearly when the error is above this threshold. Huber’s loss is convex, therefore the analysis of the previous section guarantees that ensembles will get better all the time. However, many other losses that follow this rationale are nonconvex. A nice overview of these robust regression losses was recently offered by Barron (2019), who introduced a new general loss that generalises many previous losses (both convex and nonconvex). All nonconvex losses encompassed by Barron’s generalisation (as well as others, such as using a Student’s  $t$  likelihood) share the property of being convex when the error  $(\hat{y} - y)^2$  is smaller than a prespecified threshold  $c$ , and nonconvex when the error is larger than  $c$ . For all these nonconvex losses “bad predictions” are therefore outliers and “good” ones correspond to inliers. A few of these robust losses are displayed on the right-hand side of Figure 3, all of them were chosen in order to have a cutoff of  $c = 10$ .

#### 4.1.2 MONOTONICITY OF SMOOTH NONCONVEX LOSSES

Motivated by the previous examples, we will study a general loss that is smooth, strictly convex in some part of the prediction set  $C$  (in that part, the loss Hessian will be positive definite), and strictly concave in another part of  $C$  (in that part, the loss Hessian will be negative definite). The following result shows that ensembles will eventually keep getting better when the average prediction  $\bar{y}_\infty$  ends up in this “good” part (*i.e.*, where  $L$  is strictly convex), and will keep getting worse when it ends up in the “bad” part (*i.e.* where  $L$  is strictly concave).

**Theorem 5 (Monotonicity of smooth nonconvex losses).** *Let  $\hat{y}_1, \dots, \hat{y}_K \in C$  be nondegenerate i.i.d. random variables whose first 5 moments are finite, and  $L$  be a function with continuous and bounded partial derivatives of order up to 5, with Hessian  $\nabla^2 L = H$ . Then*

1. *If  $H(\bar{y}_\infty) \succ 0$ , then the ensemble is eventually getting better: for  $K$  large enough,*

$$\mathbb{E}[L(\bar{y}_K)] < \mathbb{E}[L(\bar{y}_{K-1})], \quad (19)$$

2. *If  $H(\bar{y}_\infty) \prec 0$ , then the ensemble is eventually getting worse: for  $K$  large enough,*

$$\mathbb{E}[L(\bar{y}_K)] > \mathbb{E}[L(\bar{y}_{K-1})]. \quad (20)$$

The proof of Theorem 5 is available in Appendix C, and essentially relies on a fourth order Taylor expansion of the loss. As we will explain in Section 4.2, low-order expansions are not sufficient to prove monotonicity.



**Does this result say something about the classification error?** As illustrated in Fig. 3, the sigmoid loss can be seen as a smooth approximation of the classification error. Thus, we could expect that the classification error will behave in the same way: ensembles whose prediction are eventually correct keep getting better, and ensembles who are eventually wrong keep getting worse. One might even contemplate the idea of an easy proof using the fact that the classification error is a limiting case of a sigmoid loss. Unfortunately, we did not manage to make such a proof work as the behavior of the classification error is considerably more subtle, and involves more complex regularity conditions than the smoothness and moment assumptions of Theorem 5.

## 4.2 Why is the classification error difficult to handle?

### 4.2.1 MONOTONICITY IS A HIGHER-ORDER PHENOMENON

Since our goal is merely to show that monotonicity eventually holds, it seems natural to hope that standard asymptotic results about ensembles could do the trick. For instance, Cannings and Samworth (2017) and Lopes (2020) give conditions under which

$$\mathbb{E}[L(\bar{y}_K)] = \mathbb{E}[L(\bar{y}_\infty)] + \frac{c}{K} + o\left(\frac{1}{K}\right), \quad (21)$$

where  $L$  is the binary classification error averaged over the test set, and  $c$  is a constant. Unfortunately, it is possible to satisfy Eq. (21) and to not be eventually monotonic (consider the sequence  $1/K + (-1)^K/K^2$ ). With such expansions, monotonicity can only be captured by inspecting higher-order terms. This also explains why a fourth-order Taylor expansion is needed in the proof of Theorem 5. This means that we will need quite precise asymptotics to tackle the classification error.

### 4.2.2 A DECEIVINGLY SIMPLE COUNTER-EXAMPLE

To understand why the classification error is much more difficult to study than smooth losses, it is interesting to look at the following (counter) example, that goes back to Condorcet (1785), and was studied thoroughly by Lam and Suen (1997).

Consider a binary classification task, where the true label is 0, and our predictions  $\hat{y}_1, \dots, \hat{y}_K$  are i.i.d. samples from a Bernoulli distribution  $\mathcal{B}(\bar{y}_\infty)$ . In that case, the classification error of any  $\hat{y} \in [0, 1]$  is simply  $L(\hat{y}) = \mathbf{1}(\hat{y} \geq 0.5)$ . The expected loss of the ensemble will therefore be  $\mathbb{E}[L(\bar{y}_K)] = \mathbb{P}(\bar{y}_K \geq 0.5)$ . It seems reasonable to conjecture that  $\mathbb{P}(\bar{y}_K \geq 0.5)$  will go to zero decreasingly when  $\bar{y}_\infty < 0.5$ , and to one increasingly when  $\bar{y}_\infty > 0.5$ , which would be consistent with the analogy to the sigmoid loss sketched above. Unfortunately, this is not the case, and  $\mathbb{P}(\bar{y}_K \geq 0.5)$  is not, in general, monotonic (even asymptotically).

Indeed, let us focus for instance on the case where the ensemble is eventually right, *i.e.*, when  $\bar{y}_\infty < 0.5$ . In that case, it can be shown that the subsequence of ensembles containing an odd number of models  $L(\bar{y}_{2K+1})$  decreases (see Shteingart et al., 2020, Theorem 2, for a proof). However, monotonicity is broken when we include even numbers of models (Lam and Suen, 1997). A heuristic explanation of this failure, already mentioned by Condorcet (1785), is that even numbers of predictions allow ties, while odd numbers do not. If we alter slightly the problem to always allow ties by allowing jurors to abstain from predicting,

then monotonicity eventually holds, as we discuss in the next section and in Appendix E. Probst and Boulesteix (2018, Section 3.1.1) discuss this problem of ties as well, and possible tie-breakers.

This counterexample shows that an example as benign as a Bernoulli distribution is problematic for the classification error, and that a more complex mathematical machinery than the one of the previous proofs is required. We will develop such a machinery in the next section, based on the following simple reformulation of the expected classification error.

#### 4.2.3 A USEFUL FORMULATION OF THE EXPECTED CLASSIFICATION ERROR

We consider a classification problem with  $n_{\text{Cl}}$  classes. The true label, seen as a 1-hot encoding, is  $y \in \{0, 1\}^{n_{\text{Cl}}}$ , and a prediction is a  $n_{\text{Cl}}$ -dimensional vector  $\hat{y} = (\hat{y}^{(c)})_{c \leq n_{\text{Cl}}} \in \mathbb{R}^{n_{\text{Cl}}}$  of scores (that can be, for instance, estimated probabilities of each class, or a discrete one-hot class prediction). Without loss of generality, we assume that the true class is the first one ( $y = (1, 0, \dots, 0)$ ). The classification error is then equal to one whenever the score of the true class is smaller than the maximum score:

$$L(\hat{y}) = \begin{cases} 1 & \text{if } \hat{y}^{(1)} \leq \max_{c \neq 1} (\hat{y}^{(c)}) \\ 0 & \text{if } \hat{y}^{(1)} > \max_{c \neq 1} (\hat{y}^{(c)}) \end{cases} \quad (22)$$

The expected loss of an ensemble will thus be  $\mathbb{E}[L(\bar{y}_K)] = \mathbb{P}(\bar{y}_K^{(1)} \leq \max_{c \neq 1} \bar{y}_K^{(c)})$ , which is not a particularly appealing expression due to the presence of the maximum. However, it can be rewritten as a simple tail probability the following way:

**Lemma 6 (accuracy as margin tail probability).** *For all  $k \in \{1, \dots, K\}$ , let  $X_k = (\hat{y}_k^{(1)} - \hat{y}_k^{(2)}, \dots, \hat{y}_k^{(1)} - \hat{y}_k^{(n_{\text{Cl}})}) \in \mathbb{R}^{n_{\text{Cl}}-1}$  and  $\bar{X}_K = (\bar{y}_K^{(1)} - \bar{y}_K^{(2)}, \dots, \bar{y}_K^{(1)} - \bar{y}_K^{(n_{\text{Cl}})}) \in \mathbb{R}^{n_{\text{Cl}}-1}$ . Then,*

$$\mathbb{E}[L(\bar{y}_K)] = 1 - \mathbb{P}(\bar{X}_K > 0) = \mathbb{P}(-\bar{X}_K \geq 0), \quad (23)$$

where the inequality inside the probability is to be taken entry-wise.

**Proof**

$$\begin{aligned} \mathbb{E}[L(\bar{y}_K)] &= 1 - \mathbb{P}\left(\bar{y}_K^{(1)} > \max_{c \neq 1} \bar{y}_K^{(c)}\right) \\ &= 1 - \mathbb{P}\left(\bar{y}_K^{(1)} > \bar{y}_K^{(2)}, \dots, \bar{y}_K^{(1)} > \bar{y}_K^{(n_{\text{Cl}})}\right) \\ &= 1 - \mathbb{P}(\bar{X}_K > 0), \end{aligned}$$

by definition of  $\bar{X}_K$ . ■

The vector  $X_k \in \mathbb{R}^{n_{\text{Cl}}-1}$  corresponds to a *margin*: it stores the differences between the score of the correct label and the scores of the incorrect ones. We will discuss this further in Section 4.4.

Now that we have written the classification error as a simple tail probability, it seems natural to attack it using standard probabilistic asymptotics. As mentioned before, low-order asymptotics are not sufficient, and we will leverage “strong” large deviation theorems, that can be viewed as higher-order large deviations. This analysis is carried out in the next section, which is largely independent from the rest of the paper. We note that it may be of independent interest, beyond ensembles.

### 4.3 Monotonicity of tail probabilities

In this section, we present two results on the monotonicity of tail probabilities of empirical means. We state these results in a general form, since they may be of broader interest. Namely, we consider a sequence of i.i.d. random variables  $X_1, \dots, X_n$  living in  $\mathbb{R}^D$  with finite expectation  $\mu \in \mathbb{R}^D$ . As we hinted before, and will detail in Section 4.4, in the context of ensembles,  $X_1, \dots, X_n$  will be margins (the difference between the score of the correct class and the score of the incorrect ones). For now, we see them as generic random variables to keep this section independent from the rest of the paper.

Our goal is to find under which conditions  $\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$  is eventually strictly decreasing, for a given  $\varepsilon > 0$ . We know that  $\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$  converges to zero without additional assumptions (this is a consequence of the weak law of large numbers). We will first deal with the univariate case (Theorem 7), for which more intuitive regularity conditions are available than for the general, multivariate setting (Theorem 8). In the next section, we will be able to readily apply these results to the sequence of predictions  $\hat{y}_k$  to obtain an analogous of Theorem 5 for the classification error.

**Theorem 7 (monotonicity of tail probabilities, univariate).** *Let  $X_1, \dots, X_n$  be i.i.d. random variables with finite expectation  $\mu$ , and let  $\varepsilon > 0$ . Assume furthermore that*

- (1)  $\mathbb{E}[e^{tX_1}] < +\infty$  for all  $t \in \mathbb{R}$ ;
- (2)  $\mathbb{P}(X_1 > \mu + \varepsilon) > 0$ ;
- (3)  $X_1$  is absolutely continuous with respect to the Lebesgue measure;

or, alternatively to (3),

(3bis)  $X_1$  is a lattice random variable and  $\mathbb{P}(X_1 = \mu + \varepsilon) > 0$ .

Then,  $\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$  and  $\mathbb{P}(\bar{X}_n > \mu + \varepsilon)$  are both strictly decreasing for all  $n$  large enough.

The idea of the proof (which can be found in Section D of the Appendix) is to use precise asymptotic results on tail probabilities obtained by Petrov (1965) (see also Bahadur and Ranga Rao, 1960). These results are sometimes referred to as “sharp” or “strong” large deviation theorems (because they are more precise than standard large deviations results like Cramer’s celebrated theorem).

We want to emphasize that Theorem 7, while intuitive, may not hold if  $X_1$  does not meet the assumptions outlined in the statement of Theorem 7. As a first example, let us assume that  $X_1$  has a  $\alpha$ -stable distribution with  $\alpha < 1$ . Then one can show that  $\mathbb{P}(\bar{X}_n > \mu + \varepsilon)$  is actually *increasing* (see Section E of the Appendix). Of course, taking  $\alpha < 1$  in this example prevents  $X_1$  from having finite expectation. Another intriguing example, sketched in the previous subsection, is when  $X_1 \sim \mathcal{B}(\mu)$  with  $\mu \in (0, 1)$ . In that case, one can also derive the exact distribution of  $X_1 + \dots + X_n$ , a binomial  $\text{Bin}(n, \mu)$ . Thus  $\mathbb{P}(\bar{X}_n > \mu + \varepsilon)$  can be written as a sum of binomial coefficients, whose behavior is notoriously difficult to fully comprehend (see e.g. Lam and Suen (1997), who show non-monotonicity in some cases). Here it is assumption (3bis) which does not hold: there is no mass at  $\mu + \varepsilon$ . As we illustrate in Appendix E, adding a small mass at  $\mu + \varepsilon$  restores monotonicity. Coming back to the counter example of Section 4.2, which corresponds to the case  $\mu + \varepsilon = 0.5$ , this means

that allowing models to make no decision (i.e. outputting  $\hat{y} = 0.5$  with nonzero probability) leads to monotonicity in the classical Condorcet setting. Intuitively, this can be interpreted as follows: allowing blank votes make ties possible both when there is an odd or an even number of voters. It is quite interesting to see here the classical regularity conditions of Bahadur and Ranga Rao (1960) and Petrov (1965) explain neatly the mathematical origins of the problem of even voters in Condorcet's theorem.

Now we present a multivariate extension of Theorem 7. Note that, from now on, when  $u$  and  $v$  are both vectors,  $u \leq v$  means that  $u_i \leq v_i$  for all indices.

**Theorem 8 (Monotonicity of tails, multivariate).** *Let  $X_1, \dots, X_n \in \mathbb{R}^D$  be i.i.d. random vectors distributed as  $X$  with finite mean  $\mu$ , moment-generating function  $\phi : t \mapsto \mathbb{E}[\exp(\langle t, X_1 \rangle)]$ , and cumulant-generating function  $\varphi = \log \phi$ . Assume that*

- (1)  *$X$  has strictly positive density  $\rho$  with respect to the Lebesgue measure on  $\mathbb{R}^D$ ;*
- (2) *there exists  $\alpha > 0$  such that  $\phi$  is differentiable and bounded on  $U_\alpha$ , the open ball of radius  $\alpha$ ;*
- (3) *there exists constants  $0 < \lambda_{\min} < \lambda_{\max} < +\infty$  such that  $\text{Spec}(\nabla^2 \varphi) \subseteq [\lambda_{\min}, \lambda_{\max}]$ .*

*Then, for any  $\varepsilon \in \mathbb{R}^D$  such that  $\varepsilon > 0$  and*

$$\min_{1 \leq i \leq D} \frac{\varepsilon_i}{\|\varepsilon\|} < \sqrt{\frac{\lambda_{\min}}{\lambda_{\max}}}, \quad (24)$$

*it holds that  $\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$  is decreasing for  $n$  large enough.*

Theorem 8 is proved in Appendix D.2. In some sense, it generalises Theorem 7, since dealing with multivariate summands (which will allow us to treat the multiclass classification error in the next section). However, this comes at the cost of more stringent assumptions on the distribution of  $X$ , and in particular the local behavior of its cumulative distribution function. Moreover, while Theorem 7 was also applicable to some lattice random variables, this multivariate Theorem is restricted to absolutely continuous distributions. In particular, this means that  $\mathbb{P}(\bar{X}_n > \mu + \varepsilon) = \mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$  is also eventually decreasing.

**Remark 9 (Monotonicity of tails under Gaussian assumption).** *Let us first notice that the result of Theorem 8 is trivially true under the assumption that  $X_1, \dots, X_n$  are i.i.d.  $\mathcal{N}(0, \Sigma)$ . Indeed, in this simplified setting,*

$$\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon) = \mathbb{P}(\mathcal{N}(0, \Sigma) \geq \sqrt{n}\varepsilon),$$

*which is decreasing since we integrate with respect to a positive density.*

*Nevertheless, let us check the assumptions of Theorem 8 in the multivariate Gaussian setting in order to show that they are not too stringent. First, the density of the multivariate Gaussian does not cancel, ensuring that Assumption (1) is met. Second, it is known that*

$$\forall t \in \mathbb{R}^D, \quad \phi_X(t) = \exp\left(\mu^\top t + \frac{1}{2}t^\top \Sigma t\right),$$

which is bounded and differentiable on any open ball of finite radius, yielding Assumption (2). Third, we observe that  $\nabla^2 \phi = \Sigma$ , a positive-definite matrix. Henceforth Assumption (3) is satisfied.

Then remains the question of the existence of  $\varepsilon \in \mathbb{R}^D$  such that  $\varepsilon > 0$  and Eq. (24) holds—after all, this condition could be impossible to meet, even in such a simple scenario. We see that it is not the case: taking  $\lambda_{\min} = \lambda_{\max} = 1$  gives rise to the condition  $\min_i \varepsilon_i^2 < \|\varepsilon\|^2$ , which is true for any  $\varepsilon \in \mathbb{R}^D$  provided that  $D \geq 2$ . To see this, one can reason by contradiction and suppose that  $\min_i \varepsilon_i \geq \|\varepsilon\|$ , meaning that for all  $i \in [D]$ ,  $\varepsilon_i^2 \geq \|\varepsilon\|^2$ . Summing these inequalities returns  $\|\varepsilon\|^2 \geq D \|\varepsilon\|^2$ , a contradiction. A similar proof shows that any ratio  $\lambda_{\min}/\lambda_{\max} \geq 1/D$  does not actually put a constraint on  $\varepsilon$ .

When the  $\lambda_{\min}/\lambda_{\max}$  is smaller, the condition over  $\varepsilon$  becomes more stringent, but the set of  $\varepsilon$  remains non-empty. Indeed Eq. (24) one can always take one of the coordinates of  $\varepsilon$  to be arbitrarily small with respect to the others, automatically satisfying the condition. For instance, one can consider  $\varepsilon$  such that  $\varepsilon_1 = \sqrt{\lambda_{\min}}$  and all other coordinates equal to  $\sqrt{\lambda_{\max}}$ . Moreover, we note that Eq. (24) is a cone condition, meaning that all (positive) multiples of an  $\varepsilon$  satisfying the condition satisfies it itself. In definitive, Theorem 8 holds true for any multivariate Gaussian distribution, with a set of  $\varepsilon$  being more and more constrained, but never empty.

#### 4.4 Ensembles and the classification error

We are finally in position to deal with the classification error. We remind the notations introduced in Section 4.2.3. We deal with  $n_{\text{Cl}}$  classes, and the true label, seen as a 1-hot encoding, is  $y \in \{0, 1\}^{n_{\text{Cl}}}$ . A prediction is a  $n_{\text{Cl}}$ -dimensional vector  $\hat{y} = (\hat{y}^{(c)})_{c \leq n_{\text{Cl}}} \in \mathbb{R}^{n_{\text{Cl}}}$  of scores. Without loss of generality, we assume that the true class is the first one:  $y = (1, 0, \dots, 0)$ . Our goal is to use the results of the previous section together with the following representation, obtained in Lemma 6:

$$\mathbb{E}[L(\bar{y}_K)] = 1 - \mathbb{P}(\bar{X}_K > 0) = \mathbb{P}(-\bar{X}_K \geq 0), \quad (25)$$

where  $X_k = (\hat{y}_k^{(1)} - \hat{y}_k^{(2)}, \dots, \hat{y}_k^{(1)} - \hat{y}_k^{(n_{\text{Cl}})}) \in \mathbb{R}^{n_{\text{Cl}}-1}$ .

In the binary case, when  $n_{\text{Cl}} = 2$ ,  $X_k = \hat{y}_k^{(1)} - \hat{y}_k^{(2)}$  just corresponds to the *margin*, that is, the difference between the score of the correct label and the score of the incorrect one. In the multiclass case,  $X_k$  is a  $(n_{\text{Cl}} - 1)$ -dimensional *vector of margins* obtained for each possible “1 versus  $c$ ” classification problems, for  $c \in \{2, \dots, n_{\text{Cl}}\}$ . Following the seminal work of Bartlett et al. (1998), margins have been used a lot to study the theory of ensembles (see, e.g., Biggs et al., 2022, and references therein).

Before we state our theorem, we have to translate to the ensemble framework the regularity conditions of the two theorems of Section 4.3. We will provide two sets of assumptions that will allow us to distinguish between “good” ensembles (that eventually improve) and “bad” ones (that eventually worsen). In the rest of this section, we assume that  $\hat{y}_1, \dots, \hat{y}_k \in \mathbb{R}^{n_{\text{Cl}}}$  are i.i.d. random variables with finite mean  $\bar{y}_\infty$ .

**Assumption 1 (Correct prediction).** Let  $\varepsilon = \mathbb{E}[X_1] = \bar{y}_\infty^{(1)} 1_{n_{\text{Cl}}-1} - \bar{y}_\infty^{(-1)}$ , where we set  $u^{(-1)} = (u_2, \dots, u_p)^\top$  for any vector of  $\mathbb{R}^p$ . There are either  $n_{\text{Cl}} = 2$  classes and

the regularity conditions of Theorem 7 are verified, or there are  $n_{Cl} > 2$  classes and the regularity conditions of Theorem 8 are verified, for  $-X_1, \dots, -X_K$  and  $\varepsilon$ .

Assumption 1 implies in particular that  $\varepsilon = \mathbb{E}[X_1] > 0$  which is equivalent to the fact that the expected margin  $\bar{y}_\infty^{(1)} - \max_{c \neq 1} \bar{y}_\infty^{(c)}$  is positive. This assumption of positive margin is often used to study ensembles (see, e.g., Breiman, 2001). It implies, in particular, that the asymptotic prediction is correct, that is,  $L(\bar{y}_\infty) = 0$ .

**Assumption 2 (Completely incorrect prediction).** Let  $\varepsilon = -\mathbb{E}[X_1] = \bar{y}_\infty^{(-1)} - \bar{y}_\infty^{(1)}$ . There are either  $n_{Cl} = 2$  classes and the regularity conditions of Theorem 7 are verified, or there are  $n_{Cl} > 2$  classes and the regularity conditions of Theorem 8 are verified, for  $X_1, \dots, X_K$  and  $\varepsilon$ .

Assumption 2 implies in particular that  $\varepsilon = -\mathbb{E}[X_1] > 0$ , which means that all margins are negative, i.e., that the model asymptotically gives the lowest score to the true class. Because of this, we call this an assumption of completely incorrect prediction. This implies in particular that  $L(\bar{y}_\infty) = 1$ .

**Theorem 10 (Monotonicity of the classification error).** Let  $\hat{y}_1, \dots, \hat{y}_K \in \mathbb{R}^{n_{Cl}}$  be i.i.d. random variables. Then

1. If the prediction is asymptotically correct (Assumption 1), then the ensemble is eventually getting better: for  $K$  large enough,

$$\mathbb{E}[L(\bar{y}_K)] < \mathbb{E}[L(\bar{y}_{K-1})], \quad (26)$$

2. If the prediction is asymptotically completely incorrect (Assumption 2), then the ensemble is eventually getting worse: for  $K$  large enough,

$$\mathbb{E}[L(\bar{y}_K)] > \mathbb{E}[L(\bar{y}_{K-1})]. \quad (27)$$

**Proof** We begin by assuming that the prediction is correct (Assumption 1). Then, Theorem 7 (when  $n_{Cl} = 2$ ) or Theorem 8 (when  $n_{Cl} > 2$ ) ensures that  $\mathbb{P}(-\bar{X}_n \geq -\mathbb{E}[X_1] + \varepsilon) = \mathbb{P}(-\bar{X}_n \geq 0)$  decreases for  $K$  large enough. Using Lemma 6, this implies that  $L(\bar{y}_K)$  eventually decreases.

We now deal with the case of an incorrect prediction (Assumption 2). Again, Theorem 7 or Theorem 8 ensures that  $\mathbb{P}(\bar{X}_n > \mathbb{E}[X_1] + \varepsilon) = \mathbb{P}(\bar{X}_n > 0)$  eventually decreases. In turn,  $L(\bar{y}_K)$  eventually increases.  $\blacksquare$

## 5. Experiments

In this section, we present simple illustrations of our results on real data. The large-scale empirical study of Probst and Boulesteix (2018) can also be seen as a compelling illustration of our results in the context of random forests, as their conclusions are in line with ours.




Figure 4: Evolution of the sigmoid loss of a dropout ensemble on the dermatology data set as the number of models grows (mean and standard deviation over 500 repetitions). (Left) The sigmoid loss averaged over the whole test appears to be decreasing. (Middle) The loss averaged over images for which the asymptotic prediction  $\bar{y}_\infty$  is correct is getting better. (Right) The loss averaged over images for which the asymptotic prediction  $\bar{y}_\infty$  is incorrect is getting worse. The behaviours of these middle and right panels are explained by our theory. Notice that the three  $y$ -axes have different scales.

### 5.1 Classification of skin lesions with convolutional networks

We start by detailing the motivating example of Section 1.1. We want to predict whether a dermoscopic image corresponds to a benign keratosis-like lesion or a melanoma. We use the DermaMNIST (Yang et al., 2023) data set, based on the HAM10000 collection (Tschandl et al., 2018), and retain only the classes “benign keratosis” and “melanoma”, because they are roughly balanced within DermaMNIST (the whole dataset is very unbalanced) and notoriously difficult to distinguish (Grant-Kels et al., 1999). The training/validation/test split is the same as the one from Yang et al. (2023), and consists of 1,548 color images of resolution  $28 \times 28$  for training, 221 similar images for validation, and 443 test images.

We use a simple LeNet-like convolutional network (LeCun et al., 1998) whose fully connected layers are regularised with a dropout rate of 50% (Srivastava et al., 2014). At test time, our base predictions  $\hat{y}_1, \dots, \hat{y}_K$  are the outputs of the trained network with different random dropout masks. The ensemble prediction  $\bar{y}_K$  is called Monte Carlo model averaging by Srivastava et al. (2014, Section 7.5) and was popularised in particular by Gal and Ghahramani (2016), who called it *Monte Carlo dropout*.

As we saw in the introduction (Figure 2), the test cross-entropy is a decreasing function of  $K$ . This is explained by our theory, that indicates that, the average cross-entropy should be decreasing *for all test images*, because of Theorem 2 and the convexity of the cross-entropy. On the other hand, Figure 1 illustrated that the accuracy of the model was improving for images well-classified by the asymptotic model  $\bar{y}_\infty$ , and getting worse for the other points. This is again consistent with our theory, illustrating the dichotomy of Theorem 10.

One result that was not illustrated in the introduction was the theorem for smooth nonconvex losses (Theorem 5). To fill this gap, we look at the evolution of the sigmoid loss (see Figure 3) as the dropout ensemble grows. The result, displayed in Figure 4, tells a story




Figure 5: Evolution of the accuracy of a dropout ensemble on the 3-class dermatology data set as the number of models grows (mean and standard deviation over 500 repetitions). (*Left*) The accuracy averaged over the whole test set has a non-monotonic behaviour. (*Centre-left*) The accuracy averaged over images for which the prediction is asymptotically correct is improving. (*Centre-Right*) The accuracy averaged over images for which the prediction is asymptotically incorrect (yet the correct label is not the one with the smallest predicted probability) is degrading. (*Right*) The accuracy averaged over images for which the prediction is asymptotically completely incorrect (*i.e.* the correct label has the smallest predicted probability) is degrading. Theorem 10 provides a theoretical explanation of the behaviour of centre-left and right panels.

similar to the one of the classification error: the loss averaged over images well-classified by  $\bar{y}_\infty$  is going down, and going up when averaged over wrongly classified images, consistently with Theorem 5.

To also illustrate our results in the multiclass setting, we consider a version of the same dataset with an additional class, corresponding to basal-cell carcinoma, which is the most common type of skin cancer. We present our results in Figure 5. Here, while the assumptions of Theorem 10 are difficult to check in practice, the conclusion of the theorem seems to hold: the accuracy increases for correct predictions and decreases for completely incorrect predictions. For predictions that are neither, the accuracy is empirically decreasing. We conjecture that some extension of Theorem 10 should explain this phenomenon.

## 5.2 Wisdom of crowds for movie rating predictions

As an example of prediction of a continuous value, we now turn to an experiment closer to social sciences than machine learning. The goal is to predict the rating of an upcoming movie, using the wisdom of crowds principle, based on data collected by Simoiu et al. (2019).

Several people were asked to come up with a prediction for the future “audience score” rating of an upcoming movie on the website [www.rottentomatoes.com](http://www.rottentomatoes.com). The full dataset contains 20 movies, that were rated (out of 100) several months before their release, by around 450 users (the minimum number of ratings is 447 and the maximum 478).

For each movie and for each  $K \in \{1, \dots, 50\}$ , we sample randomly (with replacement)  $K$  of its ratings  $\hat{y}_1, \dots, \hat{y}_K \in [1, 100]$  to build an ensemble prediction  $\bar{y}_K \in [1, 100]$ . We also have a ground truth rating  $y \in [1, 100]$  for each movie, that was collected on Rotten Tomatoes after it came out. Our goal is to study if the ensemble prediction of a crowd of  $K$



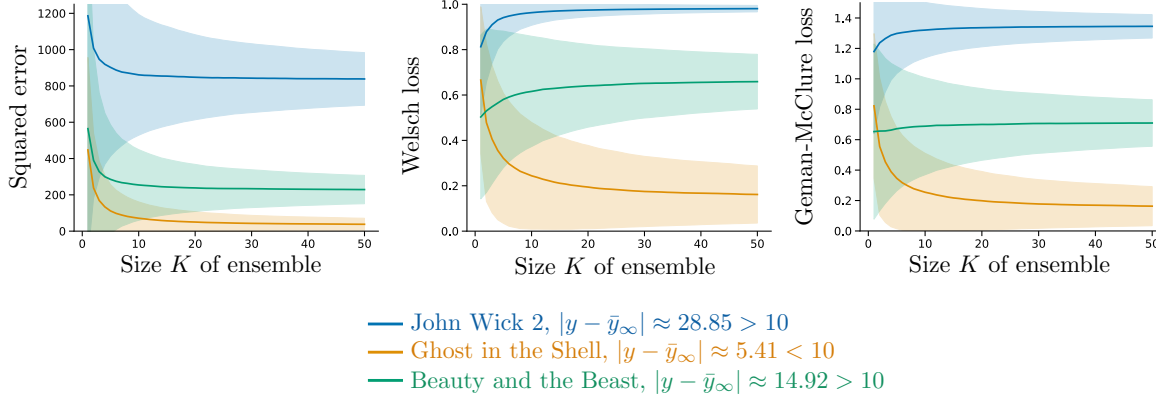


Figure 6: Evolution of the error at predicting the ratings of three movies as the size of the crowd grows (mean and standard deviation over 10,000 repetitions). *(Left)* The squared error is decreasing for all three movies. Since this is a convex loss, this is in line with Theorem 2. *(Middle and Right)* The Welsh and Geman-McClure losses decreases for Ghost in the Shell, but increases for the other two movies. This is in line with Theorem 5, which implies that, for these two losses, the loss should be eventually decreasing when  $|y - \bar{y}_\infty| < 10$ , and increasing when  $|y - \bar{y}_\infty| > 10$ .

is more accurate than the one of  $K + 1$ . To this end, we look in Figure 6 at three losses for predicting a continuous value: the squared error, the Welsh loss, and the Geman-McClure loss (for more details on these two losses, see Figure 3).

The squared error is convex, and consequently decreases, as predicted by Theorem 2. The two other losses, on the other hand, are non-convex. Their hyperparameters were chosen such that both losses are strictly convex when the error is small ( $|y - \bar{y}_\infty| < 10$ ), and strictly concave when  $|y - \bar{y}_\infty| > 10$ , as in Fig. 3. Since these two losses are additionally smooth (*i.e.*, infinitely differentiable, Barron, 2019), we can use Theorem 5 to predict that they will go up for John Wick 2 and Beauty and the Beast (because the errors are larger than 10) and go down for Ghost in the Shell (because the error is smaller than 10).

In Appendix G, we show similar curves for three more movies. It is also possible to produce such curves for all 20 movies using a Python notebook available at <https://github.com/pamattei/Getting-Better-Ensembles>.

## 6. Conclusion

Our initial question was “*Is it always true that an ensemble of  $K + 1$  models performs better than an ensemble of  $K$  models?*” When the loss is convex and the ordering of the ensemble does not matter, Theorems 2 and 4 provide a clear and positive answer to this question. When the loss is non-convex, however, a more nuanced picture must be drawn: ensembles will get better in some settings, and worse in others. Namely, ensembles will

keep getting better (respectively worse) when their asymptotic prediction falls in a locally convex (respectively concave) spot of the loss.

## 6.1 Practical guidelines

Our work leads us to strengthen the practical guidelines provided by Probst and Boulesteix (2018, Section 5.2): *in general, ensembles should be as large as computationally feasible*. Indeed, we showed that this is always true for convex losses. Furthermore, even for nonconvex losses, it is likely that, when averaged over the whole dataset, the loss will be decreasing, since it is reasonable to expect to see more “good” points than “bad” points.

An additional empirical guideline is related to evaluation of ensemble techniques. When a new technique is proposed, we recommend to report both convex and nonconvex losses when it is possible. Indeed, since *any* ensemble will get better for a convex loss, showing such improvements will hardly be useful to assess the quality of an ensemble technique. Adding other nonconvex metrics can paint a more detailed picture. Concretely, when evaluating ensembles of classifiers, we recommend to always show, at the very least, the cross-entropy (and/or the Brier score) and the classification error. This guideline is already followed quite often (e.g. by Lakshminarayanan et al., 2017, Figure 2)

## 6.2 Future work

Interesting future works include considering other aggregation functions than the mean, or looking at the whole distribution of  $L(\bar{y}_k)$  and not just its mean. In particular, it would be interesting to see if the variance of the loss, called *algorithmic variance* by Lopes (2019), is decreasing or not.

All our results were at the individual level, and can be generalised to losses averaged over datasets using the fact that averages of decreasing sequences are decreasing. However, this reasoning will not work for metrics that are not simple averages. For instance, it would be interesting to study the area under the receiver operating characteristic curve or the median squared and errors. Probst and Boulesteix (2018) noticed non-monotonic behaviours for these three losses.

Finally, it would be interesting to assess whether or not ensembling can improve model calibration. This question was investigated by researchers working on scoring rules. In particular, it was shown that ensembles of calibrated models can be uncalibrated (Hora, 2004; Ranjan and Gneiting, 2010; Gneiting and Ranjan, 2013). It would be interesting to study under which conditions calibration may be monotonically improving.

## Acknowledgments

This work was supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002. DG also acknowledges the support of ANR through project NIM-ML (ANR-21-CE23-0005-01) and of EU Horizon 2020 project AI4Media (contract no. 951911). Parts of this work were done as DG was employed at Université Côte

d’Azur. The authors thank Raphaël Razafindralambo for spotting an error in an earlier version of this paper.

## References

- S. L. Althaus. *Collective preferences in democratic politics: Opinion surveys and the will of the people*. Cambridge University Press, 2003.
- F. R. Bach. Bolasso: Model Consistent Lasso Estimation through the Bootstrap. In *Proceedings of the 25th International Conference on Machine Learning*, pages 33–40, 2008.
- R. R. Bahadur and R. Ranga Rao. On deviations of the sample mean. *The Annals of Mathematical Statistics*, 31(4):1015–1027, 1960.
- P. Baldi and P. Sadowski. The dropout learning algorithm. *Artificial Intelligence*, 210: 78–122, 2014.
- O. E. Barndorff-Nielsen. *Information and exponential families*. Wiley Series in Probability and Statistics. Wiley, 1978.
- J. T. Barron. A general and adaptive robust loss function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4331–4339, 2019.
- P. Bartlett, Y. Freund, W. S. Lee, and R. E. Schapire. Boosting the margin: A new explanation for the effectiveness of voting methods. *The Annals of Statistics*, 26(5): 1651–1686, 1998.
- P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- J. M. Bates and C. W. J. Granger. The combination of forecasts. *Journal of the Operational Research Society*, 20(4):451–468, 1969.
- D. Berend and L. Sapir. Monotonicity in condorcet jury theorem. *Social Choice and Welfare*, 24(1):83–92, 2005.
- M. Berman, A. R. Triki, and M. B. Blaschko. The Lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4413–4421, 2018.
- P. J. Bickel and K. A. Doksum. *Mathematical statistics: basic ideas and selected topics, volume I (Second edition)*. Chapman and Hall/CRC, 2015.
- F. Biggs, V. Zantedeschi, and B. Guedj. On margins and generalisation for voting classifiers. *Advances in Neural Information Processing Systems*, 35:9713–9726, 2022.
- L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

- L. D. Brown, T. T. Cai, and A. DasGupta. Interval estimation for a binomial proportion. *Statistical science*, 16(2):101–133, 2001.
- P. Bühlmann and B. Yu. Analyzing bagging. *The Annals of Statistics*, 30(4):927–961, 2002.
- Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. In *International Conference on Learning Representations*, 2016.
- T. I. Cannings. Random projections: Data perturbation for classification problems. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(1):1–15, 2021.
- T. I. Cannings and R. J. Samworth. Random-projection ensemble classification. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(4):959–1035, 2017.
- A. Casa, L. Scrucca, and G. Menardi. Better than the best? Answers via model ensemble in density-based clustering. *Advances in Data Analysis and Classification*, 15:599–623, 2021.
- J. Chu-Carroll, K. Czuba, J. Prager, and A. Ittycheriah. In question answering, two heads are better than one. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 24–31, 2003.
- R. T. Clemen. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5(4):559–583, 1989.
- N. Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, 1785. Reprinted by Cambridge University Press in 2014.
- O. Dekel and O. Shamir. Vox Populi: Collecting High-Quality Labels from a Crowd. In *COLT*, 2009.
- L. Devroye, L. Györfi, and G. Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- P. Diaconis. Recent progress on de Finetti’s notions of exchangeability. *Bayesian Statistics*, 3:111–125, 1988.
- K. L. Dougherty and J. Edward. Odd or even: Assembly size and majority rule. *The Journal of politics*, 71(2):733–747, 2009.
- L. Elkin and R. Pettigrew. *Opinion Pooling*. Cambridge University Press, 2025.
- W. Feller. *An introduction to probability theory and its applications, vol. II*. John Wiley & Sons, 1971.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.

- J. H. Friedman and P. Hall. On bagging and nonlinear estimation. *Journal of Statistical Planning and Inference*, 137(3):669–683, 2007.
- Y. Gal and Z. Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016.
- F. Galton. Vox populi. *Nature*, 75:450–451, 1907.
- C. Genest and J. V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–135, 1986.
- E. I. George. Bayesian model averaging: A tutorial: Comment. *Statistical Science*, 14(4):409–412, 1999.
- T. Gneiting and A. E. Raftery. Weather forecasting with ensemble methods. *Science*, 310(5746):248–249, 2005.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.
- T. Gneiting and R. Ranjan. Combining predictive distributions. *Electronic Journal of Statistics*, 7:1747–1782, 2013.
- J. M. Grant-Kels, E. T. Bason, and C. M. Grin. The misdiagnosis of malignant melanoma. *Journal of the American Academy of Dermatology*, 40(4):539–548, 1999.
- L. Grinsztajn, E. Oyallon, and G. Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? In *Thirty-sixth Conference on Neural Information Processing Systems, Datasets and Benchmarks Track*, 2022.
- P. Hall. *The bootstrap and Edgeworth expansion*. Springer Series in Statistics. Springer Science & Business Media, 1992.
- L. K. Hansen and P. Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):993–1001, 1990.
- T. K. Ho. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):832–844, 1998.
- J. A. Hoeting, D. Madigan, A. E. Raftery, and C. T. Volinsky. Bayesian model averaging: a tutorial (with comments by M. Clyde, D. Draper, and E. I. George, and a rejoinder by the authors. *Statistical Science*, 14(4):382–417, 1999.
- S. C. Hora. Probability judgments for continuous quantities: Linear combinations and calibration. *Management Science*, 50(5):597–604, 2004.
- J. Hron, A. Matthews, and Z. Ghahramani. Variational Bayesian dropout: pitfalls and fixes. In *International Conference on Machine Learning*, pages 2019–2028. PMLR, 2018.

- P. J. Huber. Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35:73–101, 1964.
- C. Joutard. Multidimensional strong large deviation results. *Metrika*, 80(6):663–683, 2017.
- A. Koriati. When are two heads better than one and why? *Science*, 336(6079):360–362, 2012.
- A. Krogh and J. Vedelsby. Neural network ensembles, cross validation, and active learning. *Advances in Neural Information Processing Systems*, 7, 1994.
- L. I. Kuncheva. *Combining pattern classifiers: methods and algorithms*. John Wiley & Sons, 2014.
- A. Laan, G. Madirolas, and G. G. De Polavieja. Rescuing collective wisdom when the average group opinion is wrong. *Frontiers in Robotics and AI*, 4(56):1–21, 2017.
- K. K. Ladha. Condorcet’s jury theorem in light of de Finetti’s theorem: Majority-rule voting with correlated votes. *Social Choice and Welfare*, 10:69–85, 1993.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing Systems*, 30, 2017.
- L. Lam and S. 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.
- K. G. Larsen. Bagging is an optimal PAC learner. In *Conference on Learning Theory*, pages 450–458. PMLR, 2023.
- O. Laurent, A. Lafage, E. Tartaglione, G. Daniel, J.-M. Martinez, A. Bursuc, and G. Franchi. Packed ensembles for efficient uncertainty estimation. In *International Conference on Learning Representations*, 2023.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- M. E. Lopes. Estimating the algorithmic variance of randomized ensembles via the bootstrap. *The Annals of Statistics*, 47(2):1088–1112, 2019.
- M. E. Lopes. Estimating a sharp convergence bound for randomized ensembles. *Journal of Statistical Planning and Inference*, 204:35–44, 2020.
- X. Ma, H. Huang, Y. Wang, S. Romano, S. Erfani, and J. Bailey. Normalized loss functions for deep learning with noisy labels. In *International Conference on Machine Learning*, pages 6543–6553. PMLR, 2020.
- C. F. Manski. When consensus choice dominates individualism: Jensen’s inequality and collective decisions under uncertainty. *Quantitative Economics*, 1(1):187–202, 2010.

- C. F. Manski. Interpreting point predictions: Some logical issues. *Foundations and Trends® in Accounting*, 10(2-4):238–261, 2016.
- A. W. Marshall and F. Proschan. An inequality for convex functions involving majorization. *Journal of Mathematical Analysis and Applications*, 12(1):87–90, 1965.
- A. W. Marshall, I. Olkin, and B. Arnold. *Inequalities: Theory of Majorization and Its Applications*. Springer Series in Statistics. Springer, 2011.
- H. Masnadi-Shirazi and N. Vasconcelos. On the design of loss functions for classification: theory, robustness to outliers, and SavageBoost. *Advances in Neural Information Processing Systems*, 21, 2008.
- H. Masnadi-Shirazi, V. Mahadevan, and N. Vasconcelos. On the design of robust classifiers for computer vision. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 779–786. IEEE, 2010.
- L. Mason, J. Baxter, P. Bartlett, and M. Frean. Boosting algorithms as gradient descent. *Advances in Neural Information Processing Systems*, 12, 1999.
- P.-A. Mattei and J. Frellsen. Uphill roads to variational tightness: Monotonicity and Monte Carlo objectives. *arXiv preprint arXiv:2201.10989*, 2022.
- P. McCullagh. *Tensor methods in statistics (Second edition)*. Dover Books on Mathematics, 2018.
- S. K. McNees. The uses and abuses of ‘consensus’ forecasts. *Journal of Forecasting*, 11(8): 703–710, 1992.
- K. Nikodem. On strongly convex functions and related classes of functions. In *Handbook of functional equations*, pages 365–405. Springer, 2014.
- H. Noh, T. You, J. Mun, and B. Han. Regularizing deep neural networks by noise: Its interpretation and optimization. *Advances in Neural Information Processing Systems*, 30, 2017.
- L. A. Ortega, R. Cabañas, and A. Masegosa. Diversity and generalization in neural network ensembles. In *International Conference on Artificial Intelligence and Statistics*, pages 11720–11743. PMLR, 2022.
- G. Parmigiani and L. Inoue. *Decision theory: Principles and approaches*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2009.
- B. Peleg and S. Zamir. Extending the Condorcet jury theorem to a general dependent jury. *Social Choice and Welfare*, 39:91–125, 2012.
- M. P. Perrone. *Improving regression estimation: Averaging methods for variance reduction with extensions to general convex measure optimization*. PhD thesis, Brown University, 1993.

- V. V. Petrov. On the probabilities of large deviations for sums of independent random variables. *Theory of Probability & Its Applications*, 10(2):287–298, 1965.
- D. Prelec, H. S. Seung, and J. McCoy. A solution to the single-question crowd wisdom problem. *Nature*, 541(7638):532–535, 2017.
- P. Probst and A.-L. Boulesteix. To tune or not to tune the number of trees in random forest. *Journal of Machine Learning Research*, 18:1–18, 2018.
- R. Ranjan and T. Gneiting. Combining probability forecasts. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 72(1):71–91, 2010.
- C. E. Rasmussen and C. K. Williams. *Gaussian processes for machine learning*, volume 1. Springer, 2006.
- A. J. Riesselman, J. B. Ingraham, and D. S. Marks. Deep generative models of genetic variation capture the effects of mutations. *Nature methods*, 15(10):816–822, 2018.
- L. Rokach. Ensemble-based classifiers. *Artificial Intelligence Review*, 33:1–39, 2010.
- N. Russell, T. B. Murphy, and A. E. Raftery. Bayesian model averaging in model-based clustering and density estimation. *arXiv preprint arXiv:1506.09035*, 2015.
- E. Scornet and G. Hooker. Theory of random forests: A review. *HAL preprint hal-05006431*, 2025.
- D. Serre. *Matrices. Theory and Applications (Second edition)*. Graduate Texts in Mathematics. Springer, 2010.
- S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge University Press, 2014.
- H. Shteingart, E. Marom, I. Itkin, G. Shabat, M. Kolomenkin, M. Salhov, and L. Katzir. Majority Voting and the Condorcet’s Jury Theorem. *arXiv preprint arXiv:2002.03153*, 2020.
- C. Simoiu, C. Sumanth, A. Mysore, and S. Goel. Studying the “wisdom of crowds” at scale. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 171–179, 2019.
- C. G. Small. *Expansions and asymptotics for statistics*. Number 115 in Monographs on Statistics and Applied Probability. Chapman and Hall/CRC, 2010.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- M. Stone. The opinion pool. *The Annals of Mathematical Statistics*, pages 1339–1342, 1961.



- L. Struski, M. Mazur, P. Batorski, P. Spurek, and J. Tabor. Bounding Evidence and Estimating Log-Likelihood in VAE. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 5036–5051. PMLR, 25–27 Apr 2023.
- R. Sun, A. Ramé, C. Masson, N. Thome, and M. Cord. Towards efficient feature sharing in mimo architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2697–2701, 2022.
- J. Surowiecki. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday / Anchor, 2004.
- R. Theisen, H. Kim, Y. Yang, L. Hodgkinson, and M. W. Mahoney. When are ensembles really effective? *Advances in Neural Information Processing Systems*, 36:15015–15026, 2023.
- P. Tschandl, C. Rosendahl, and H. Kittler. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018.
- T. Viering and M. Loog. The shape of learning curves: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7799–7819, 2023.
- T. Viering, A. Mey, and M. Loog. Open problem: Monotonicity of learning. In *Conference on Learning Theory*, pages 3198–3201. PMLR, 2019.
- V. D. Wild, S. Ghalebikesabi, D. Sejdinovic, and J. Knoblauch. A rigorous link between deep ensembles and (variational) bayesian methods. *Advances in Neural Information Processing Systems*, 36:39782–39811, 2023.
- R. L. Winkler. Probabilistic prediction: Some experimental results. *Journal of the American Statistical Association*, 66(336):675–685, 1971.
- J. Yang, R. Shi, D. Wei, Z. Liu, L. Zhao, B. Ke, H. Pfister, and B. Ni. MedMNIST v2-A large-scale lightweight benchmark for 2D and 3D biomedical image classification. *Scientific Data*, 10(1):41, 2023.

## Appendix A. Proof of strict improvements under strong convexity

We prove here Theorem 4. Strong convexity leads to the following strengthening of Jensen's inequality (see, *e.g.*, Nikodem, 2014, Theorem 2).

**Theorem 11 (Jensen improvement).** *If  $L : C \rightarrow \mathbb{R}$  is  $\mu$ -strongly convex, then, for all  $x_1, \dots, x_K$ , and*

$$L\left(\frac{1}{K} \sum_{k=1}^K x_k\right) \leq \frac{1}{K} \sum_{k=1}^K L(x_k) - \mu \frac{1}{K} \sum_{k=1}^K \left\| x_k - \frac{1}{K} \sum_{j=1}^K x_j \right\|^2. \quad (28)$$

Our goal is to combine this result with our basic identity Eq. (8) to prove Theorem 4. We use Theorem 11 with  $x_j = \frac{1}{K-1} \sum_{k \neq j} \hat{y}_k$ . Using Eq. (8) and the exchangeability assumption, we find that

$$\mathbb{E} \left[ L\left(\frac{1}{K} \sum_{k=1}^K \hat{y}_k\right) \right] \leq \mathbb{E} \left[ L\left(\frac{1}{K-1} \sum_{k=1}^{K-1} \hat{y}_k\right) \right] - \mu \mathbb{E} \left[ \left\| \frac{1}{K-1} \sum_{k=1}^{K-1} \hat{y}_k - \frac{1}{K} \sum_{k=1}^K \hat{y}_k \right\|^2 \right]. \quad (29)$$

We can then compute the expectation on the right-hand side

$$\begin{aligned} \mathbb{E} \left[ \left\| \frac{1}{K-1} \sum_{k=1}^{K-1} \hat{y}_k - \frac{1}{K} \sum_{k=1}^K \hat{y}_k \right\|^2 \right] &= \frac{1}{K^2(K-1)^2} \mathbb{E} \left[ \left\| K \sum_{k=1}^{K-1} \hat{y}_k - (K-1) \left( \hat{y}_K + \sum_{k=1}^{K-1} \hat{y}_k \right) \right\|^2 \right] \\ &= \frac{1}{K^2(K-1)^2} \mathbb{E} \left[ \left\| \sum_{k=1}^{K-1} \hat{y}_k - (K-1) \hat{y}_K \right\|^2 \right] \\ &= \frac{1}{K^2(K-1)^2} \text{tr} \left( \text{Cov} \left( \sum_{k=1}^{K-1} \hat{y}_k - (K-1) \hat{y}_K \right) \right). \end{aligned} \quad (30)$$

We can now use Bienaymé's identity and exchangeability to compute the trace of the covariance:

$$\begin{aligned} \text{tr} \left( \text{Cov} \left( \sum_{k=1}^{K-1} \hat{y}_k - (K-1) \hat{y}_K \right) \right) &= (K-1) \text{tr}(\Sigma) + (K-1)(K-2) \text{tr}(\text{Cov}(\hat{y}_1, \hat{y}_2)) \\ &\quad + (K-1)^2 \text{tr}(\Sigma) - 2 \text{tr} \left( \text{Cov} \left( \sum_{k=1}^{K-1} \hat{y}_k, (K-1) \hat{y}_K \right) \right), \end{aligned} \quad (31)$$

finally leading to

$$\text{tr} \left( \text{Cov} \left( \sum_{k=1}^{K-1} \hat{y}_k - (K-1) \hat{y}_K \right) \right) = K(K-1) \text{tr}(\Sigma) (1 - \rho). \quad (32)$$

Plugging this expression into Equation (30) gives the desired bound. ■

## Appendix B. More details on nonconvex losses

### B.1 More details on the spherical score

Consider the setting of Section 4.1. The spherical score is defined as

$$L_1(\hat{y}) = -\frac{\hat{y}}{\sqrt{\hat{y}^2 + (1 - \hat{y})^2}}, \quad (33)$$

when the true label is  $y = 1$ . When the true label is 0, as in Section 4.1, it is defined as

$$L_0(\hat{y}) = -\frac{1 - \hat{y}}{\sqrt{\hat{y}^2 + (1 - \hat{y})^2}}. \quad (34)$$

The second derivative of  $L_0$  is  $\hat{y} \mapsto (-4\hat{y}^2 + \hat{y} + 1)/(2(\hat{y} - 1)x + 1)^{5/2}$ , whose unique root in  $[0, 1]$  is  $1/8 + \sqrt{17}/8 \approx 0.64$ . Therefore, as seen in Figure 3,  $L_0$  will be strictly convex on  $[0, 1/8 + \sqrt{17}/8)$  and strictly concave on  $(1/8 + \sqrt{17}/8, 1]$ .

### B.2 More details on the tangent loss

The tangent loss (Masnadi-Shirazi et al., 2010) is a loss function for binary classification. There are two labels,  $y = 1$  and  $y = -1$ , and the output of our model is a score  $\hat{y} \in \mathbb{R}$  that indicates its confidence in predicting that the label is 1. When the true label is  $y \in \{-1, 1\}$ , it is defined as

$$L_y(\hat{y}) = (2\arctan(y\hat{y}) - 1)^2. \quad (35)$$

This function is convex around its minimum of one, but nonconvex when predictions are overconfident. Indeed, when  $y = 1$ , the function is locally convex on  $[y_-, y_+]$  with  $y_- \approx -0.51$  and  $y_+ \approx 1.26$ , and locally concave elsewhere.

## Appendix C. Proof of the theorem for smooth losses (Theorem 5)

We only treat the case  $H(\hat{y}_\infty) \succ 0$ , since the case  $H(\hat{y}_\infty) \prec 0$  follows by replacing  $L$  by  $-L$ . We will base our proof on the following lemma, which generalises a result of Small (2010, Section 4.3), and is obtained by Taylor-expanding the loss up to the fourth order. In this context, such a technique is often called the delta method for moments (see, e.g., Small, 2010; Bickel and Doksum, 2015). The proof of the lemma is slightly tedious, and is provided after the main theorem.

**Lemma 12 (fourth order delta method).** *Let  $X_1, \dots, X_n \in \mathbb{R}^d$  be a sequence of non-degenerate i.i.d. random variables, whose first 5 moments are finite, with mean  $\mu$  and covariance matrix  $\Sigma \succeq 0$ . Let  $L$  be a function with continuous and bounded partial derivatives of order up to 5, with Hessian matrix  $H$ . Then, there exists a constant  $\alpha \in \mathbb{R}$ , that depends on the first 3 moments of  $X$  and the partial derivatives of the loss of order up to 4, such that, when  $K$  goes to infinity*

$$\mathbb{E}[L(\bar{X}_K)] = L(\mu) + \frac{\text{tr}(H(\mu)\Sigma)}{2K} + \frac{\alpha}{K^2} + \mathcal{O}\left(\frac{1}{K^{5/2}}\right). \quad (36)$$

*Proof of Theorem 5.* Applying Lemma 12 successively to  $\bar{y}_K$  and  $\bar{y}_{K+1}$  and taking the difference between the two terms, we obtain

$$\mathbb{E}[L(\bar{y}_K)] - \mathbb{E}[L(\bar{y}_{K+1})] = \frac{\text{tr}(H(\hat{y}_\infty)\Sigma)}{2} \left( \frac{1}{K} - \frac{1}{K+1} \right) + \alpha \left( \frac{1}{K^2} - \frac{1}{(K+1)^2} \right) + \mathcal{O} \left( \frac{1}{K^{5/2}} \right) \quad (37)$$

$$= \frac{\text{tr}(H(\hat{y}_\infty)\Sigma)}{2} \frac{1}{K(K+1)} + \alpha \frac{1+2K}{K^2(K+1)^2} + \mathcal{O} \left( \frac{1}{K^{5/2}} \right). \quad (38)$$

Therefore,

$$K(K+1) (\mathbb{E}[L(\bar{y}_K)] - \mathbb{E}[L(\bar{y}_{K+1})]) = \frac{\text{tr}(H(\hat{y}_\infty)\Sigma)}{2} + \alpha \frac{1+2K}{K(K+1)} + \mathcal{O} \left( \frac{1}{K^{1/2}} \right), \quad (39)$$

which implies that  $K(K+1) (\mathbb{E}[L(\bar{y}_K)] - \mathbb{E}[L(\bar{y}_{K+1})])$  will converge to  $\text{tr}(H(\hat{y}_\infty)\Sigma)/2$  when  $K$  goes to infinity. Therefore, for  $K$  large enough, the sign of  $\mathbb{E}[L(\bar{y}_K)] - \mathbb{E}[L(\bar{y}_{K+1})]$  is constant and exactly equal to the sign of  $\text{tr}(H(\hat{y}_\infty)\Sigma)$ .

Since  $H(\hat{y}_\infty) \succ 0$ ,  $\Sigma \succeq 0$ , and  $\Sigma \neq 0$  because we assumed that the random variable is nondegenerate, a standard property of symmetric matrices (see, *e.g.*, Serre, 2010, Proposition 6.1) ensures that  $H(\hat{y}_\infty)\Sigma$  is diagonalisable, with only real nonnegative eigenvalues (at least one of them being positive). This implies that  $\text{tr}(H(\hat{y}_\infty)\Sigma) > 0$ , and, in turn, that for  $K$  large enough,  $\mathbb{E}[L(\bar{y}_{K+1})] < \mathbb{E}[L(\bar{y}_K)]$ .  $\blacksquare$

### C.1 Proof of the fourth order delta method (Lemma 12)

Lemma 12 was proven in the univariate case by Small (2010, Section 4.3) and we extend it here to the multivariate case. We will see it as a consequence of the multivariate version of delta method presented by Bickel and Doksum (2015).

Under the assumptions of Lemma 12, Theorem 5.3.2 of Bickel and Doksum (2015) implies

$$\mathbb{E}[L(\bar{X}_K)] = L(\mu) + \sum_{r=2}^4 \frac{1}{r!} \sum_{j_1, j_2, \dots, j_r=1}^d \frac{\partial^r L}{\partial x_{j_1} \dots \partial x_{j_r}}(\mu) \mathbb{E}[(\bar{X}_{Kj_1} - \mu_{j_1}) \dots (\bar{X}_{Kj_r} - \mu_{j_r})] + \mathcal{O}(K^{-5/2}). \quad (40)$$

Let us look at the quadratic term. We have

$$\begin{aligned} \frac{1}{2} \sum_{j_1, j_2=1}^d \frac{\partial^2 L}{\partial x_{j_1} \partial x_{j_2}}(\mu) \mathbb{E}[(\bar{X}_{Kj_1} - \mu_{j_1})(\bar{X}_{Kj_2} - \mu_{j_2})] &= \frac{1}{2} \mathbb{E}[(\bar{X}_K - \mu)^T H(\mu)(\bar{X}_K - \mu)] \quad (41) \\ &= \frac{1}{2} \text{tr}(H(\mu) \text{Cov}(\bar{X}_K)) = \frac{\text{tr}(H(\mu)\Sigma)}{2K}. \quad (42) \end{aligned}$$

Let us now look at the third and fourth order terms. We need to assess how  $K$  affects the multivariate moments of the empirical mean. Let us define first the multivariate moments of  $X$  as, for  $r \in \{3, 4\}$  and  $j_1, \dots, j_r \in \{1, \dots, d\}$

$$\mu_{j_1, \dots, j_r} = \mathbb{E}[(X_{j_1} - \mu_{j_1}) \dots (X_{j_r} - \mu_{j_r})], \quad (43)$$

The moments of  $\bar{X}_K$  can now be expressed in terms of moments of  $X$ . Following for instance Hall (1992, p. 54) or McCullagh (2018, Section 2.3), we have

$$\mathbb{E}[(\bar{X}_{Kj_1} - \mu_{j_1})(\bar{X}_{Kj_2} - \mu_{j_2})(\bar{X}_{Kj_3} - \mu_{j_3})] = \frac{1}{K^2} \mu_{j_1, j_2, j_3}, \quad (44)$$

and

$$\begin{aligned} \mathbb{E}[(\bar{X}_{Kj_1} - \mu_{j_1})(\bar{X}_{Kj_2} - \mu_{j_2})(\bar{X}_{Kj_3} - \mu_{j_3})(\bar{X}_{Kj_4} - \mu_{j_4})] = \\ \frac{\mu_{j_1, j_2} \mu_{j_3, j_4} + \mu_{j_1, j_3} \mu_{j_2, j_4} + \mu_{j_1, j_4} \mu_{j_2, j_3}}{K^2} + \mathcal{O}\left(\frac{1}{K^3}\right), \end{aligned} \quad (45)$$

where the constant in the  $\mathcal{O}$  is a multivariate cumulant. Plugging these expressions into Equation (40) finally gives

$$\mathbb{E}[L(\bar{X}_K)] = L(\mu) + \frac{\text{tr}(H(\mu)\Sigma)}{2K} + \frac{\alpha}{K^2} + \mathcal{O}\left(\frac{1}{K^{5/2}}\right), \quad (46)$$

with

$$\begin{aligned} \alpha = \frac{1}{3!} \sum_{j_1, j_2, j_3=1}^d \frac{\partial^3 L}{\partial x_{j_1} \partial x_{j_2} \partial x_{j_3}}(\mu) \mu_{j_1, j_2, j_3} \\ + \frac{1}{4!} \sum_{j_1, j_2, j_3, j_4=1}^d \frac{\partial^4 L}{\partial x_{j_1} \partial x_{j_2} \partial x_{j_3} \partial x_{j_4}}(\mu) (\mu_{j_1, j_2} \mu_{j_3, j_4} + \mu_{j_1, j_3} \mu_{j_2, j_4} + \mu_{j_1, j_4} \mu_{j_2, j_3}). \end{aligned} \quad (47)$$

One can check that, in the univariate case ( $d = 1$ ), our expression for  $\alpha$  is identical to the one given by Small (2010, Section 4.3).  $\blacksquare$

## Appendix D. Omitted proofs for Section 4.3

### D.1 Proof of Theorem 7

For all  $i \geq 1$ , let us set  $Y_i := X_i - \mu$ . Note that, by definition, the  $Y_i$ s are i.i.d. and centered. We are going to apply two theorems of Petrov (1965) to the random variables  $Y_i$ s. Let us recall briefly Petrov's results here. We start with some notation:  $V(x)$  denotes the cumulative distribution function of  $X_1$ ,  $\mathfrak{U}$  denotes the points of increase of  $V$  and  $A$  its supremum,  $\mathcal{B}^+$  denotes the set of non-negative values  $h$  such that  $\int_0^{+\infty} e^{hx} dV(x) < +\infty$  and  $B$  its supremum. Further, assuming  $B > 0$ , for any  $0 < h < B$ , we define

$$R(h) = \int_{-\infty}^{+\infty} e^{hx} dV(x), \quad m(h) = \frac{1}{R(h)} \int_{-\infty}^{+\infty} x e^{hx} dV(x) \quad \text{and} \quad \sigma(h)^2 = \frac{dm(h)}{dh}.$$

Finally, define  $A_0 = \lim_{h \rightarrow B^-} m(h)$ . We are now able to state the results:

**Theorem 13 (Theorem 3 and Theorem 6 from Petrov (1965)).** *Let  $Y_1, Y_2, \dots$  be a sequence of i.i.d. random variables.*

- Assume that the  $Y_i$ s are non-lattice such that  $B > 0$ . Let  $c$  be any constant satisfying the condition  $\mathbb{E}[Y_1] < c < A_0$ . Then

$$\mathbb{P}(Y_1 + \dots + Y_n \geq nc) = \frac{\exp(n \log R(h) - nhc)}{h\sigma(h)\sqrt{2\pi n}}(1 + o(1)).$$

Here  $h$  is the unique real root of the equation  $m(h) = c$ .

- Assume that the  $Y_i$  are lattice distributed with values in  $a + NH$ ,  $N \in \mathbb{Z}$ . Assume that  $\mathbb{E}[Y_1] > -\infty$ ,  $A_0 < +\infty$ , and  $B > 0$ . If  $nx \in na + \mathbb{Z}H$ , then

$$\mathbb{P}(Y_1 + \dots + Y_n = nx) = \frac{H \exp(n \log R(h) - nhx)}{\sigma(h)\sqrt{2\pi n}} \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right),$$

and

$$\mathbb{P}(Y_1 + \dots + Y_n \geq nx) \leq \frac{H \exp(n \log R(h) - nhx)}{\sigma(h)\sqrt{2\pi n}(1 - e^{-Hh})} \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right)$$

as  $n \rightarrow +\infty$  uniformly in  $x \in [\mathbb{E}[Y_1] + \varepsilon, A_0 - \varepsilon]$ , with  $\varepsilon > 0$  arbitrary. Here  $h$  is the unique real root of the equation  $m(h) = x$ .

Theorem 13 gives us a control of the tail probability, which is precise enough to show monotonicity. There are two cases to consider: the *non-lattice* case (assuming (1), (2), and (3)) and the *lattice* case (assuming (1), (2), and (3bis)). We denote  $p_n = \mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$ , for  $\varepsilon$  satisfying the assumptions of the theorem.

**Non-lattice case.** Under assumption (3), we know that the  $Y_i$ s are non-lattice random variables. Let us check carefully the assumptions of Petrov's result: Under assumption (1), we know that  $B = +\infty$ . Under assumption (2), we know that  $\mathbb{P}(Y_1 > \varepsilon) > 0$ . In the language of Petrov (1965), it implies that there is a growth point of  $Y_1$  greater than  $\varepsilon$ . Thus  $A$  is greater than  $\varepsilon$ . Since we are in the case  $B = +\infty$ , we know that  $A_0 = A$ , and in particular  $A_0 > \varepsilon > 0$ . Thus we can apply the first part of Theorem 13 to the  $Y_i$ s with  $c = \varepsilon$ . We obtain

$$\forall n \in \mathbb{N}^*, \quad p_n = \frac{\exp(n \log R(h) - nh\varepsilon)}{h\sigma(h)\sqrt{2\pi n}}(1 + u_n), \quad (48)$$

where  $u_n = o(1)$  as  $n \rightarrow +\infty$ . From Eq. (48), we deduce that, for all  $n \geq 1$ ,

$$\frac{p_{n+1}}{p_n} = \rho \sqrt{\frac{n}{n+1}} \frac{1 + u_{n+1}}{1 + u_n}, \quad (49)$$

where  $\rho := \exp(\log R(h) - h\varepsilon)$ . By the first lemma in Petrov (1965),  $m$  is strictly increasing and continuous in  $(0, +\infty)$ . Thus, by definition of  $h$ , the mapping  $\alpha : t \mapsto \log R(t) - t\varepsilon$  is decreasing for  $t < h$  and increasing for  $t > h$ , attaining a (strict) minimum in  $t = h$ . Since  $\alpha(0) = 0$ , we deduce that  $\alpha(h) < 0$ , which in turn implies  $\rho < 1$  in Eq. (49). Since both  $\sqrt{n/(n+1)}$  and  $(1 + u_{n+1})/(1 + u_n)$  are arbitrarily close to 1 for large  $n$ , we deduce the result in the case of a non-lattice distribution.

**Lattice case.** As in the previous paragraph, under assumption (1) and (2),  $B = +\infty$  and  $A_0 > \varepsilon > 0$ . Moreover, assumption (3bis) guarantees that  $n\varepsilon$  takes values in the lattice. Similarly to Eq. (48), we obtain

$$\forall n \in \mathbb{N}^*, \quad p_n = \frac{H \exp(n \log R(h) - nh\varepsilon)}{\sigma(h) \sqrt{2\pi n} (1 - e^{-Hh})} (1 + u_n), \quad (50)$$

where  $u_n = o(1)$ . From Eq. (50), we deduce that, for all  $n \geq 1$ ,

$$\frac{p_{n+1}}{p_n} = \rho \sqrt{\frac{n}{n+1}} \frac{1 + u_{n+1}}{1 + u_n}, \quad (51)$$

where  $\rho$  is as before. We conclude in the same fashion as in the first part of the proof and obtain monotonicity for  $\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$ .

Let us now deal with the strict inequality case. Decomposing  $\mathbb{P}(\bar{X}_n > \mu + \varepsilon)$  as the difference between  $\mathbb{P}(\bar{X}_n \geq \mu + \varepsilon)$  and  $\mathbb{P}(\bar{X}_n = \mu + \varepsilon)$ , and using again Theorem 13, we obtain

$$\mathbb{P}(\bar{X}_n > \mu + \varepsilon) = \frac{H \exp(n \log R(h) - nh\varepsilon)}{h \sigma(h) \sqrt{2\pi n}} \left( \frac{1}{1 - e^{-Hh}} - 1 \right) (1 + u_n), \quad (52)$$

where  $u_n = o(1)$ . Since  $1/(1 - e^{-Hh}) - 1 > 0$ , the previous reasoning applies, and  $\mathbb{P}(\bar{X}_n > \mu + \varepsilon)$  is also eventually monotonic.  $\blacksquare$

**Remark 14.** *It is possible to relax Assumption (1): a careful reading of Petrov (1965) reveals that one only needs a non-degenerate interval  $[0, B)$  on which the moment generating function of  $X_1$  is well-defined, with  $B > 0$ . But in that event,  $\varepsilon$  may need to be taken “small enough.” If Assumption (2) is not satisfied, then simple counter-examples such as a Dirac at  $\mu$  break the statement. Assumption (3bis) cannot be relaxed, as will be seen in the second example of the following discussion.*

## D.2 Proof of Theorem 8

Let us set  $a := \mu + \varepsilon$ . We first notice that the assumptions of Lemma 15 (which is stated and proved in Section F.1) are satisfied, thus there exists  $\tau \in U_a$  such that  $\nabla \varphi(\tau) = a$  and  $\tau_{(1)} > 0$ . We now prove a stronger result than Theorem 8, namely, for  $n$  large enough,

$$\mathbb{P}(\bar{X}_n \geq a) = \frac{\exp(-n(\tau^\top a - \varphi(\tau)))}{(2\pi n)^{D/2} \prod_{k=1}^D \tau_k \det(\nabla^2 \varphi(\tau))^{1/2}} (1 + o(1)). \quad (53)$$

The result follows from this last display, since  $\tau^\top a - \varphi(a) > 0$ . Indeed,  $t \mapsto t^\top a - \varphi(t)$  is a strictly concave function, with maximum achieved at  $t = \tau$  and value 0 at 0. All that is left to do is to prove Eq. (53). We use Theorem 1 in Joutard (2017). More precisely, we apply Theorem 1 to  $Z_n = \bar{X}_n$ , with rate  $b_n = n$  and  $H = 0$ . We have several assumptions to check.

First, in Joutard (2017)’s notation, the moment-generating function of  $b_n Z_n$  is given by  $\phi_n(t) = \phi(t)^n$ , since the  $X_i$ s are i.i.d. Simply by taking the log, we obtain that the normalized cumulant generating function of  $b_n Z_n$  is given by  $\varphi_n(t) = \log \phi(t)$ . By assumption,  $\phi$

is defined and differentiable on  $U_\alpha$ . Since  $\phi$  is positive by construction, we have shown the existence of a differentiable function  $\varphi$  such that  $\lim_n \varphi_n(t) = \varphi(t)$  for all  $t \in U_\alpha$ .

We now proceed to show that assumption A1-3 are satisfied. We start with A1, asking that  $\varphi_n$  is an holomorphic function on  $D_\alpha^p$ , where  $D_\alpha := \{z \in \mathbb{C}\}$ . We notice that, since  $\varphi_n = \varphi$ , we only have to show the claim for  $\varphi$ . Let us fix all variables except the  $i$ th. That is, let us consider the mapping  $\varphi_i : s \mapsto \varphi(t_1, \dots, t_{i-1}, s, t_{i+1}, \dots, t_p)$  for some  $t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_p \in \mathbb{R}$ . By our assumption on  $\varphi$ ,  $\varphi_i$  is well-defined on the segment  $(-\alpha, \alpha)$ , and therefore  $\varphi_i$  as a function of the complex variable is well-defined and holomorphic on the strip  $\{z = a + ib, a \in (-\alpha, \alpha), b \in \mathbb{R}\}$ . In particular,  $\varphi_i$  is holomorphic on the disk  $D_\alpha$ . Hartogs' theorem tells us that, since  $\varphi : D_\alpha^p \rightarrow \mathbb{C}$  is a continuous function. Finally, Osgood's Lemma guarantees that  $\varphi$  is holomorphic on  $D_\alpha^p$ . That it is bounded is straightforward from our assumption.

Let us now turn to A2. Again, since  $\varphi_n = \varphi$ , the first part of the assumption is trivial to check (one can take  $H(t) = 0$ ). In order to prove that  $\nabla^2 \varphi(\tau)$  is positive definite, we apply exponential tilting, also called by Petrov (1965) the *Cramér method*. More precisely, let us define the random vector  $Z$  with density with respect to Lebesgue measure on  $\mathbb{R}^p$  defined as

$$\frac{1}{\phi(\tau)} \exp(\langle \tau, x \rangle) \rho(x),$$

where, as we recall,  $\rho$  is the density of  $X$ . Then one can readily check that  $\nabla^2 \varphi(\tau) = \text{Cov}(Z)$ . In particular,  $\nabla^2 \varphi(\tau)$  is positive semi-definite. Since, by assumption,  $X$  has non-zero density everywhere, the same is true for  $Z$  by construction. In particular,  $Z$  is not supported in an hyperplane. Thus its covariance matrix has no zero eigenvalue and we can conclude.

All that is left to do is to prove that A3 is satisfied. We first notice that, according to the discussion at the beginning of this proof, for any  $t \in \mathbb{R}$ ,

$$\left| \frac{\phi_n(\tau + \mathbf{i}t)}{\phi_n(\tau)} \right| = \left| \frac{\phi(\tau + \mathbf{i}t)}{\phi(\tau)} \right|. \quad (54)$$

Thus showing that  $\left| \frac{\phi(\tau + \mathbf{i}t)}{\phi(\tau)} \right| < 1$  will guarantee that Eq. (54) is  $o(1/b_n^{p/2})$ . But one can recognize that

$$\frac{\phi(\tau + \mathbf{i}t)}{\phi(\tau)} = \mathbb{E}[\exp(\mathbf{i}\langle t, Z \rangle)],$$

which is the characteristic function of the random vector  $Z$  defined in the previous paragraph. Again, since  $X$  has non-zero density everywhere, the same is true for  $Z$  and Lemma 17 (which is proved in Section F.2) guarantees that  $|\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)]| < 1$  for any  $t \neq 0$ .  $\blacksquare$

## Appendix E. When monotonicity does not hold

In this section, we give two examples in which the assumptions of Theorem 7 are not satisfied and  $(p_n)_{n \geq 1}$  is not monotonous. We set  $S_n := X_1 + \dots + X_n$ .



**Example 1.** Let us assume that  $X_1$  has a  $\alpha$ -stable distribution, with density  $f(x; \alpha, \beta, c, \mu)$ . In that case,  $S_n - n\mu = \sum_{i=1}^n 1 \cdot (X_i - \mu)$  has density

$$\frac{1}{n^{1/\alpha}} f(y/n^{1/\alpha}; \alpha, \beta, c, \mu).$$

We can directly compute  $p_n$  by a change of variables, and we obtain

$$\begin{aligned} p_n &= \mathbb{P}(S_n - n\mu > n\varepsilon) \\ &= \int_{n\varepsilon}^{+\infty} \frac{1}{n^{1/\alpha}} f(y/n^{1/\alpha}; \alpha, \beta, c, \mu) dy \\ p_n &= \int_{n^{1-1/\alpha}\varepsilon}^{+\infty} f(x; \alpha, \beta, c, \mu) dx. \end{aligned}$$

Whenever  $\alpha < 1$ ,  $(n^{1-1/\alpha}\varepsilon)_n$  is a decreasing sequence, and according to the last display  $p_n$  is increasing. This is even though the distribution is symmetric around zero. We showcase this effect in Figure 7.

**Example 2.** Let us assume that  $X_1 \sim \mathcal{B}(p)$ , and therefore  $S_n \sim \text{Bin}(n, p)$ . Thus

$$\begin{aligned} p_n &= \mathbb{P}(X_1 + \dots + X_n \geq n(p + \varepsilon)) \\ &= \sum_{k=\lceil n(p+\varepsilon) \rceil}^n \binom{n}{k} p^k (1-p)^{n-k}. \end{aligned}$$

While innocent-looking, this last display has a complicated behavior. In particular, it is not monotonic (see Figure 8). Transforming slightly this example in order to satisfy the assumptions of Theorem 7 radically changes the behavior of  $p_n$ , as predicted. One possible way to do so is to add a small mass  $\varepsilon/(2(\mu + \varepsilon) - 1)$  at  $\mu + \varepsilon$ , while setting the probability to be equal to 0 and 1 equal to  $(2\mu - 1)/(2(\mu + \varepsilon) - 1)$ . We refer to Figure 8 for an illustration.

## Appendix F. Technical results

### F.1 On the existence of $\tau$

For a subset  $A$  of  $\mathbb{R}^D$ , we denote by  $C(A)$  its convex hull, and by  $\text{int}(A)$  its interior.

**Lemma 15 (Existence of  $\tau$ ).** Let  $X \in \mathbb{R}^D$  be a random variable with logarithmic moment generating function  $\varphi$ , support  $\mathcal{S}$ , and mean  $\mu$ . We assume that  $X$  is absolutely continuous with respect to the Lebesgue measure, and that there exists  $\alpha \in ]0, \infty]$  such that  $\{t \in \mathbb{R}^D \mid \varphi(t) < \infty\} = U_\alpha$ , where  $U_\alpha$  denotes the ball of radius  $\alpha$ . Let  $\lambda_{\min}$  (resp.  $\lambda_{\max}$ ) denote the minimal (resp. maximal) eigenvalue of  $\nabla^2 \varphi$ . Then, for all  $\varepsilon \in \mathbb{R}^D$  with strictly positive entries such that  $\mu + \varepsilon \in \text{int}(C(\mathcal{S}))$  and

$$\min_{1 \leq i \leq D} \frac{\varepsilon_i}{\|\varepsilon\|} > \sqrt{2 \left( 1 - \sqrt{\frac{\lambda_{\min}}{\lambda_{\max}}} \right)}, \quad (55)$$

there exists  $\tau \in U_\alpha$  with strictly positive entries such that  $\nabla \varphi(\tau) = \mu + \varepsilon$ .




Figure 7: Evolution of  $p_n$  when  $X_1$  follows the  $\alpha$ -stable distribution. Here we considered  $\varepsilon = 0.1$ . When  $\alpha < 1$ ,  $p_n$  is increasing while it is decreasing, as expected, when  $\alpha > 1$ .




Figure 8: *Left*: evolution of  $p_n$  when  $X_1$  follows a Bernoulli distribution. Here we took  $\varepsilon = 0.1$ . *Right*: same example, with some mass added at  $\mu + \varepsilon$ . We observe that the sequence  $(p_n)_{n \geq 1}$  is decreasing for sufficiently large  $n$ , as predicted by Theorem 7.

**Remark 16.** *Assumption (55) is quite restrictive. Indeed, assuming for a moment that  $\|\varepsilon\| = 1$ , we see that the locus of points such that  $\varepsilon_i > c$  for all  $i$  is empty whenever  $c \geq 1/\sqrt{d}$ . Therefore, one needs*

$$1 - \frac{1}{2d} \leq \frac{\lambda_{\min}}{\lambda_{\max}} (\leq 1)$$

for Eq. (55) to hold for a non-empty set of points. For large  $d$ , this means requiring nearly isotropic  $\nabla^2 \varphi$ .

**Proof** Let  $\varepsilon \in \mathbb{R}^D$  satisfying our assumptions. Since  $\{t \in \mathbb{R}^D | \varphi(t) < \infty\} = U_\alpha$ ,  $\varphi$  is “steep” in the sense of Barndorff-Nielsen (1978), and his Theorem 9.2 implies then that  $\nabla \varphi(U_\alpha) = \text{int}(C(\mathcal{S}))$ . Therefore, there exists  $\tau \in U_\alpha$  such that  $\nabla \varphi(\tau) = \mu + \varepsilon$ .

All that is left to do is to show that all entries of  $\tau$  are strictly positive. Since  $\varphi$  is  $\lambda_{\max}$ -smooth, one can write

$$\frac{1}{\lambda_{\max}} \|\nabla \varphi(\tau) - \nabla \varphi(0)\|^2 \leq (\nabla \varphi(\tau) - \nabla \varphi(0))^\top (\tau - 0) = (\mu + \varepsilon - \mu)^\top \tau = \varepsilon^\top \tau. \quad (56)$$

Thus we have proved that  $\varepsilon^\top \tau \geq \frac{1}{\lambda_{\max}} \|\varepsilon\|^2$ . In particular,  $\tau$  points in the same direction as  $\varepsilon$ . Though this is not sufficient to conclude: the halfspace of directing vector  $\varepsilon$  contains points having non-positive coordinates. We note that, since  $X$  is absolutely continuous,  $\lambda_{\min} > 0$ , and  $\varphi$  is  $\lambda_{\min}$ -strongly convex. Therefore,

$$(\nabla \varphi(\tau) - \nabla \varphi(0))^\top (\tau - 0) \geq \lambda_{\min} \|\tau - 0\|^2, \quad (57)$$

that is,  $\varepsilon^\top \tau \geq \lambda_{\min} \|\tau\|^2$ . Combining both bounds, we see that

$$\varepsilon^\top \tau \geq \sqrt{\frac{1}{\lambda_{\max}} \|\varepsilon\|^2 \cdot \lambda_{\min} \|\tau\|^2},$$

from which we deduce that

$$\cos(\varepsilon, \tau) \geq \sqrt{\frac{\lambda_{\min}}{\lambda_{\max}}}. \quad (58)$$

Equivalently, the angle between  $\varepsilon$  and  $\tau$  is less than  $\theta_0 := \arccos\left(\sqrt{\frac{\lambda_{\min}}{\lambda_{\max}}}\right)$ , meaning that  $\tau$  belongs to  $\mathcal{C}$  the half-cone of direction  $\varepsilon$  and aperture  $2\theta_0$ .

Let us now prove that, under our assumption, no point of  $\mathcal{C}$  has a negative coordinate. Without loss of generality, we can assume that  $\|\varepsilon\| = \|\tau\| = 1$ . Using Eq. (58), we see that

$$\|\varepsilon - \tau\|^2 = 1 + 1 - 2\varepsilon^\top \tau \leq 2 \left(1 - \sqrt{\frac{\lambda_{\min}}{\lambda_{\max}}}\right) =: \rho^2.$$

To put it plainly,  $\tau$  belongs to  $\mathcal{B}$ , the closed ball of center  $\varepsilon$  and radius  $\rho$ . The distance from the center of  $\mathcal{B}$  to each hyperplane  $\{x_i = 0\}$  is given by  $\varepsilon_i$ . Our assumption guarantees that  $\varepsilon_i > \rho$  for all  $i \in [d]$ , thus  $\mathcal{B}$  does not intersect any of the hyperplanes  $\{x_i = 0\}$ . In particular, all point inside  $\mathcal{B}$  have positive coordinates, and since  $\tau \in \mathcal{B}$  we can conclude. ■




Figure 9: Visual help for the proof of Lemma 15. The vector  $\tau$  is contained in a half cone of axis  $\varepsilon$  and angle  $\theta_0$  (in red).

## F.2 On the characteristic function of a random vector with non-zero density

In this section, we state and prove the following result, which is a straightforward multivariate extension of Lemma 3 in Chapter XV.7 of Feller (1971).

**Lemma 17 (Characteristic function of random vector).** *Let  $X \in \mathbb{R}^p$  be a random vector with non-zero density almost everywhere. Then*

$$\forall t \neq 0, \quad |\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)]| < 1.$$

**Proof** Let us notice that the triangle inequality guarantees that  $|\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)]| \leq 1$  for any  $t \in \mathbb{R}^p$ . Therefore, one only needs to show that  $|\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)]| \neq 1$ , which we will achieve reasoning by contradiction.

First, let us suppose that there exists  $t \neq 0$  such that  $\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)] = 1$ . Then the expectation of the non-negative function  $x \mapsto 1 - \cos\langle t, x \rangle$  is exactly 0, meaning that  $\langle t, x \rangle = 0$  modulo  $2\pi$  wherever  $X$  has non-zero mass. Since  $t \neq 0$ , this implies that  $X$  is supported in the collection of hyperplanes defined by

$$\langle t, x \rangle = 2k\pi, \quad k \in \mathbb{Z},$$

which contradicts  $X$  having non-zero density everywhere.

Let us now come back to the original problem, and assume that  $|\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)]| = 1$ . Therefore there exists  $b \in \mathbb{R}$  such that  $|\mathbb{E}[\exp(\mathbf{i}\langle t, X \rangle)]| = e^{ib}$ . Since  $t \neq 0$ , there exists a vector  $c \in \mathbb{R}^p$  such that  $b = \langle t, c \rangle$ , and one can write

$$\mathbb{E}[\exp(\mathbf{i}\langle t, X - c \rangle)] = 1.$$

We recognize the characteristic function of the random vector  $Y := X - c$ , which also has non-zero density everywhere. Applying the previous reasoning to  $Y$  also brings contradiction, and we can conclude.  $\blacksquare$

## Appendix G. Three more movies

The data from Simoiu et al. (2019) deal with 20 movies, and only three were displayed in the main paper. We look here at three more movies: Logan, The Space Between Us, and Smurfs, The Lost Village. As illustrated in Figure 10, the results are again in line with our theory. An interesting feature of the Smurfs movie is that while it is in the nonconvex part of the loss, it is close to the inflexion point, since  $|y - \bar{y}_\infty| \approx 13.8$  is closer to 10 than any of the five other movies. This perhaps explains while the Welsch and Geman-McClure losses are initially decreasing, and eventually increasing.

Convergence is in general much slower when  $\bar{y}_\infty$  is close to the inflection point of 10, than when raters are either very accurate (like for The Space Between Us), or clearly missed the target (John Wick). This can be explained by the delta method expansion of Lemma 12, that implies

$$\mathbb{E}[L(\bar{y}_k)] = L(\bar{y}_\infty) + \frac{L''(\bar{y}_\infty)\text{Var}(\hat{y}_1)}{2K} + \mathcal{O}\left(\frac{1}{K^{3/2}}\right), \quad (59)$$

which means that convergence will be slower when the second derivative is small, *i.e.* near the inflexion point.

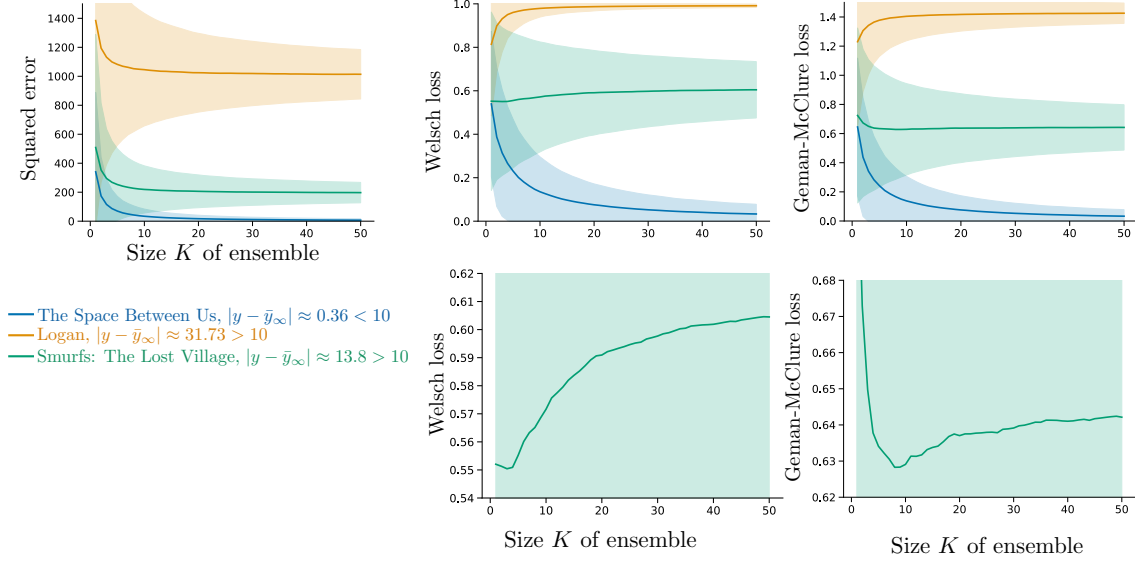


Figure 10: Evolution of the error at predicting the ratings of three movies as the size of the crowd grows (mean and standard deviation over 10,000 repetitions). *(Left)* The squared error is decreasing for all three movies, as predicted by Theorem 2. *(Middle and Right)* The Welsh and Geman-McClure losses are clearly increasing for Logan, and decreasing for The Space Between us, in line with Theorem 5. The Smurfs are not monochromatic as we expected: as displayed in the two zoom-ins of the bottom panels, the nonconvex losses are decreasing at first, then end-up increasing. This remains consistent with Theorem 5, that merely ensures that these two loss will be *eventually* monotonic.