



HAL
open science

Sequential Aggregation of Probabilistic Forecasts -Applicaton to Wind Speed Ensemble Forecasts

Michaël Zamo, Liliane Bel, Olivier Mestre

► **To cite this version:**

Michaël Zamo, Liliane Bel, Olivier Mestre. Sequential Aggregation of Probabilistic Forecasts - Applicaton to Wind Speed Ensemble Forecasts. *Journal of the Royal Statistical Society: Series C Applied Statistics*, 2021, 70 (1), pp.202-225. 10.1111/rssc.12455 . hal-03753771

HAL Id: hal-03753771

<https://agroparistech.hal.science/hal-03753771v1>

Submitted on 25 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sequential Aggregation of Probabilistic Forecasts - Application to Wind Speed Ensemble Forecasts

Michaël Zamo^{1,3}, Liliane Bel², and Olivier Mestre^{1,3}

¹Météo-France, Direction des Opérations pour la Production, 42
avenue Gaspard Coriolis, 31057 Toulouse cedex 07, France.

²Université Paris-Saclay, AgroParisTech, INRAE, UMR
MIA-Paris, 75005, Paris, France

³CNRM/GAME, Météo-France/CNRS URA 1357, Toulouse,
France

May 8, 2020

Abstract

In the field of numerical weather prediction (NWP), the probabilistic distribution of the future state of the atmosphere is sampled with Monte-Carlo-like simulations, called ensembles. These ensembles have deficiencies (such as conditional biases) that can be corrected thanks to statistical post-processing methods. Several ensembles exist and may be corrected with different statistical methods. A further step is to combine these raw or post-processed ensembles. The theory of prediction with expert advice allows us to build combination algorithms with theoretical guarantees on the forecast performance. This article adapts this theory to the case of probabilistic forecasts issued as step-wise cumulative distribution functions (CDF). The theory is applied to wind speed forecasting, by combining several raw or post-processed ensembles, considered as CDFs. The second goal of this study is to explore the use of two forecast performance criteria: the Continuous ranked probability score (CRPS) and the Jolliffe-Primo test. Comparing the results obtained with both criteria leads to reconsidering the usual way to build skillful probabilistic forecasts, based on the minimization of the CRPS. Minimizing the CRPS does not necessarily produce reliable forecasts according to the Jolliffe-Primo test. The Jolliffe-Primo test generally selects reliable forecasts, but could lead to issuing suboptimal forecasts in terms of CRPS. It is proposed to use both criterion to achieve reliable and skillful probabilistic forecasts.

1 Introduction

As a chaotic dynamical system, the atmosphere has an evolution that is intrinsically uncertain (Malardel, 2005; Holton and Hakim, 2012) and should be described in a probabilistic form. In the field of numerical weather prediction (NWP), this probabilistic form is not a probability distribution but a set of deterministic forecasts whose aim is to assess the forecast uncertainty (Leutbecher and Palmer, 2008). Such a set of deterministic forecasts is called an ensemble forecast and each individual deterministic forecast is called a member. The members are usually obtained by running the same NWP model with different initial conditions and different parametrizations of the model physics (Descamps et al., 2011). Forecast uncertainty can then be derived from the members as a probability distribution with statistical estimation techniques and considering the members are a random sample from an unknown multivariate probability distribution .

Being often biased and under-dispersed for surface parameters (Hamill and Colucci, 1998; Buizza et al., 2005), the ensemble forecast systems may be post-processed with statistical methods, called ensemble model output statistics (EMOS) to get more skillful forecast distributions (Wilson et al., 2007; Thorarinsdottir and Gneiting, 2010; Möller and Scheuerer, 2013; Zamo et al., 2014; Baran and Lerch, 2015; Taillardat et al., 2016).

Nowadays, several ensemble forecast systems are available routinely (Bougeault et al., 2010; Descamps et al., 2014). Combining, or “aggregating”, several forecasts may improve the predictive performance compared to the most skillful post-processed ensemble (Allard et al., 2012; Gneiting et al., 2013; Baudin, 2015; Baran and Lerch, 2016; Möller and Groß, 2016; Bogner et al., 2017). The theory of prediction with expert advice (Cesa-Bianchi et al., 2006; Stoltz, 2010) shows how to efficiently aggregate in real-time several forecasts based on their respective past performances. This theory studies the mathematical properties of aggregation algorithms of several forecasting systems (called “experts” in this framework), and has been applied mostly to point forecasts.

The first goal of the present work is to apply the theory of prediction with expert advice to probabilistic forecasts represented as step-wise cumulative distribution functions (CDF) with any number of steps. Two previous studies (Baudin, 2015; Thorey, 2017) used this theory for specific cases of probabilistic forecasts. In Baudin (2015), the experts are the *ordered* individual members of pooled ensembles. Each expert’s forecast is a stepwise cumulative distribution function with one step. In this case the experts are not identifiable over time although it is required by the theory. For instance, at different times, the lowest forecast value comes from a different member of a different ensemble. Thorey (2017) applies the theory of prediction with expert advice to forecasts of photovoltaic electricity production. The experts are ensemble forecasts or built from ensemble or deterministic forecasts thanks to statistical regression methods. Each expert is treated as a different deterministic forecast, even though some of them

are members of the same ensemble. Other experts are built thanks to quantile regression methods. As such, the set of forecast quantiles could be considered as a specific probabilistic expert, but each forecast quantile is again considered as a separate deterministic expert. In this study, contrary to Baudin (2015) and Thorey (2017), each expert is a whole (raw or post-processed) ensemble, and is thus actually identifiable over time. This work extends the work of Baudin (2015) in the sense that the formulae established in this article reduce to the ones in Baudin (2015) if considering step-wise CDFs with a single step. In a nutshell, the aggregated forecast is a linear combination of the CDFs of the experts. Since the aggregated forecast must be a CDF, only convex aggregation strategies are investigated: the weights are constrained to be positive and to sum up to one.

The second goal of this work is to compare two model selection approaches when dealing with probabilistic forecasting systems (Collet and Richard, 2017). The first approach is based on “reliability” (or “calibration”) and “sharpness” (Gneiting et al., 2007; Jolliffe and Stephenson, 2011). A forecasting system is reliable if the conditional probability distribution of the observation given the forecast distribution is equal to the forecast distribution. A forecasting system is sharper when, on average, it predicts a lower dispersion of the observation. According to the *sharpness-calibration paradigm* of Gneiting et al. (2007) a forecasting system should aim at providing reliable probabilistic forecasts that are the sharpest (i.e. less dispersed) possible. A practical motivation of this sharpness-calibration paradigm is that decisions based on such a forecasting system would be optimal due to the reliability of the forecast and less uncertain due to the forecast’s low dispersion, which improves its value for economical decisions (Richardson, 2001; Zhu et al., 2002; Mylne, 2002). Among several reliable forecasting systems, one should thus select the sharpest. The second approach to model selection among probabilistic forecasting systems is based on a scoring rule, such as the Continuous Ranked Probabilistic Score (CRPS, Matheson and Winkler 1976): the selected model is the one that has the best value of the scoring rule (highest or lowest value depending on the scoring rule). The two approaches to model selection do not yield equivalent forecasts, as previously mentioned in different studies (Collet and Richard, 2017; Wilks, 2018). For instance, minimizing the CRPS may lead to forecasts that are not reliable. To solve this problem, Wilks (2018) proposes to minimize the CRPS penalized with a term quantifying the unreliability of the forecast. Collet and Richard (2017) introduces a post-processing method that, under quite strong assumptions, yields reliable forecasts without degrading too much the CRPS compared to the CRPS-minimizing approach. We do not use any of these solutions here but compare the two model selection approaches and their properties on a case study. Both approaches are used to select the best forecasting system among several experts and aggregated forecasts. In the first approach, reliability is imposed by using the Jolliffe-Primo flatness test (JP test, Jolliffe and Primo (2008)). As for the second approach, forecast performance is measured with the CRPS.

In Section 2, the theory of prediction with expert advice is presented, along with notations. It is shown how this theory can be straightforwardly applied to step-wise CDFs. The CRPS and JP test are also introduced with more details and their use is further motivated. Section 3 presents the different aggregation methods investigated. Some are empirical, while others exhibit interesting theoretical properties. Section 4 describes the use-case of this study and the data it uses: four ensemble forecasts, two EMOS methods used to post-process the ensembles, and the wind speed observation. The results of the comparison of the aggregation methods are presented in Section 5. These results motivate a more theoretical comparison of the two approaches to model selection among probabilistic forecasts, in Section 6. Finally, Section 7 concludes with a summary of the results and perspectives.

2 Theoretical Framework and Performance Assessment Tools

The desired properties of forecast aggregation are two-fold. The first one is to yield an aggregated forecast that performs better than any of the forecasts that are used in the aggregation. According to the theory of prediction with expert advice (Cesa-Bianchi et al., 2006; Stoltz, 2010), some algorithms used to sequentially aggregate forecasts exhibits theoretical guarantees of performance. These guarantees state that the aggregated forecast will not perform much worse than some skillful reference forecast, called the oracle. In practice, the aggregated forecast may even outperform the oracle. This motivates using the theory of prediction with expert advice. The second desired property is to dynamically tackle changes in the forecasts' generating process (such as modification in NWP model's code). These changes may strongly affect the performance of the raw or post-processed forecasts. A good aggregation method should quickly detect changes in the performance of the individual forecasts and adapt the aggregation weights to discard the bad ones and favor the good ones. Being a sequential aggregation framework based on the recent performance of the experts, the theory of prediction with expert advice may help in reaching this second goal.

2.1 Sequential Aggregation of Step-Wise CDFs

The situation tackled by the theory of prediction with expert advice is the following: a forecaster has to forecast some parameter of interest by using only past observations of the parameter and past and current forecasts of the parameter stemming from several sources. These sources whose forecasts are to be aggregated are called "experts" in this theory. In this very general framework no assumption is made on the generating process of the observations and the experts.

More formally, at time t , before the observation $y_t \in \mathcal{Y}$ is revealed, let us suppose available the forecast of the $E \in \mathbb{N}^*$ experts and of the past observations (noted $y_{j=1,\dots,t-1}$). An expert is any means, in a very general sense (NWP model, human expertise, ...), to produce a forecast of y_t at each time t , before the observation y_t is known. The forecast of expert $e \in \{1, \dots, E\}$ at time t is noted $\hat{y}_{e;t} \in \hat{\mathcal{Y}}$, with $\hat{\mathcal{Y}}$ the value set of the forecasts. Although the prediction with expert advice has been mostly used with experts issuing point forecasts ($\hat{\mathcal{Y}} \subseteq \mathbb{R}$), let us stress that the experts yield forecasts of any type, not necessarily point forecasts. The theory straightforwardly adapts to probabilistic forecasts as will be shown below. Someone or something, called the “forecaster” in the theory, produces an aggregated forecast \hat{y}_t as a linear combination of the experts’ current forecasts

$$\hat{y}_t = \sum_{e=1}^E \omega_{e;t} \hat{y}_{e;t}, \quad (1)$$

where $\omega_{e;t} \in \Omega \subseteq \mathbb{R}$ is the aggregation weight of expert e at time t . The aggregation weights are computed using only information available at the present time, namely the past and present experts’ forecasts $\hat{y}_{e;j}$ with $e = 1, \dots, E$ and $j = 1, \dots, t$, and the past observations y_j , with $j = 1, \dots, t-1$. The aggregation is usually initialized with equal weights, that is, $\omega_{e;1} = \frac{1}{E}$, $\forall e = 1, \dots, E$. The weights $\omega_{e;t}$ can be computed with many algorithms (called aggregation methods), some of which are presented in Section 3.

When the observation y_t is revealed, the forecaster suffers a loss, quantified with a function $\ell : \hat{\mathcal{Y}} \times \mathcal{Y} \rightarrow \mathbb{R}$. The most general goal of the forecaster would be to build the best possible forecast, that is, to minimize its cumulative loss over a period of time $t = 1, \dots, T$. This minimization is not possible for all sequences $(\hat{y}_{e;t}, y_t)_{e,t}$, since it is always theoretically possible to build a sequence of observations that makes the forecaster’s cumulative loss arbitrarily high. Therefore, a more realistic goal is to build the best possible forecast relatively to the best element from some class of reference forecasts. Let us note \mathcal{C} such a class of reference forecasts, whose elements are written

$$\tilde{y}_t = \sum_{e=1}^E \tilde{\omega}_{e;t} \hat{y}_{e;t}, \quad (2)$$

In this case the weights $\tilde{\omega}_{e;t}$ may be computed by using the present and future information (that is on the whole period $t' \in \{1, \dots, T\}$). Hence the aggregated forecast delivered by the forecaster may be different from the best aggregated forecast in class \mathcal{C} (called the oracle). Two examples of class \mathcal{C} are the set of the available experts, or the set of linear combination of the experts’ forecasts with constant weighting computed with some chosen aggregation method. The computation of the oracle uses all the information for the whole period. Thus, the oracle cannot be used for real-time applications, but may be used as a reference in order to evaluate aggregation methods. The regret $R_T^{\mathcal{C}}$ of the forecaster relatively to the class \mathcal{C} is the cumulative additional loss suffered by the forecaster

who used its own aggregated forecast instead of the oracle of class \mathcal{C} :

$$R_T^{\mathcal{C}} = \sum_{t=1}^T \ell(\hat{y}_t, y_t) - \inf_{\tilde{y}_t \in \mathcal{C}} \sum_{t=1}^T \ell(\tilde{y}_t, y_t). \quad (3)$$

According to the theory of prediction with expert advice there exists aggregation methods available to the forecaster such that the regret relative to class \mathcal{C} is sub-linear in T

$$\sup R_T^{\mathcal{C}} \leq o(T), \quad (4)$$

where the supremum is taken over all possible sequences of the observation and experts. Let's point out that the regret is not lower bounded, and so may be negative for some datasets. Since the upper bound on the regret holds for *any* sequence of observation and forecasts, this sublinearity property ensures to the forecaster good forecast performances. In most cases, this property only requires that the chosen loss ℓ be convex in its first argument, and no assumption is required about the observation or the experts (Cesa-Bianchi et al., 2006). In specific cases where some information is known about the experts, such as a correlation structure between the experts' forecasts, this information may be used to design better aggregation methods (see supplement in Swinbank et al. (2016)). Adjakossa et al. (2020) uses assumptions about the experts' generating process to improve aggregation methods for point forecasts. Here only the general case is treated. Mallet et al. (2007) and Gerchinovitz et al. (2008) review many aggregation methods of expert advice, with numerical algorithms thereof. Both articles describe the case of experts issuing point forecasts ($\hat{\mathcal{Y}} \subseteq \mathbb{R}$) and real observations ($\mathcal{Y} \subseteq \mathbb{R}$), along with theoretical bounds for the L_2 -loss ($\ell(\hat{y}_t, y_t) = (\hat{y}_t - y_t)^2$), when they exist.

Applying this theory for probabilistic forecasts expressed as step-wise CDF is straightforward. The observation is supposed real-valued ($\mathcal{Y} \subseteq \mathbb{R}$). The forecast space $\hat{\mathcal{Y}}$ is the set of step-wise CDF, that is the set of piece-wise constant, non-decreasing functions taking their values in $[0; 1]$. Each expert forecast $\hat{y}_{e;t}$ is a step function with jumps of heights $p_e^{m_e}$ (called weights) at $M_e \in \mathbb{N}^*$ values $x_{e;t}^{m_e} \in \mathcal{Y}$. For instance the values $x_{e;t}^{m_e}$ may come from an ensemble forecast with M_e members. The weights are such that $p_e^{m_e} > 0$, for $m_e = 1, \dots, M_e$, and $\sum_{m_e=1}^{M_e} p_e^{m_e} = 1$. Then each expert's forecast is the step-wise CDF

$$\hat{y}_{e;t} = \sum_{m_e=1}^{M_e} p_e^{m_e} H(x - x_{e;t}^{m_e}) \quad (5)$$

where H is the Heaviside function, and $x \in \mathcal{Y}$. Without loss of generality, the $x_{e;t}^{m_e}$ are supposed sorted in ascending order for each expert and at each time t , so that $x_{e;t}^{m_e}$ may be considered as the quantile of order $\tau_e^{m_e} = \sum_{m'_e=1}^{m_e} p_e^{m'_e}$ of a random variable $\hat{Y}_{e;t} \in \mathcal{Y}$.

The aggregated forecast CDF, \hat{y}_t , is a step function at the pooled values $\{x_{e;t}^{m_e}; m_e =$

$1, \dots, M_e, e = 1, \dots, E\}$ such that the jump of height $\omega_{e;t}p_e^{m_e}$ is associated to the value $x_{e;t}^{m_e}$, thus

$$\hat{y}_t(x_{e;t}^{m_e}) = \sum_{e'=1}^E \omega_{e';t} \left(\sum_{m'=1}^{M_{e'}} p_{e'}^{m'} H(x_{e;t}^{m_e} - x_{e';t}^{m'}) \right) = \tau_{e;t}^{m_e}. \quad (6)$$

In other words, $x_{e;t}^{m_e}$ may be considered as the quantile of order $\tau_{e;t}^{m_e}$ of some random variable $\hat{Y}_t \in \mathcal{Y}$, whose computation is illustrated in Figure 1. To produce a valid step-wise CDF, the aggregation method must produce aggregation weights that are all non negative and sum up to 1 at fixed t . Thus, the aggregated forecast is a convex combination of the expert forecasts.




Figure 1 – Example of aggregation of $E = 2$ CDFs. The forecast CDFs (blue continuous line for expert $e = 1$, red dashed line for expert $e = 2$) are known only through a set of $M_1 = M_2 = 3$ values $x_e^{m_e}$ with associated jumps $p_e^{m_e} = \frac{1}{3}$. Following Equation (6), since x_2^1 is greater than x_1^1 and x_1^2 , it is the quantile of order $\tau_{2;t}^1 = \omega_{1;t}(p_1^1 + p_1^2) + \omega_{2;t}p_2^1 = \frac{2}{3}\omega_{1;t} + \frac{1}{3}\omega_{2;t}$ of \hat{y}_t . If $\omega_{1;t} = \omega_{2;t} = \frac{1}{2}$, then $\tau_{2;t}^1 = \frac{1}{2}$. In this case, the aggregated step-wise CDF is the black continuous line.

2.2 Performance Assessment Tools

Among the several attributes of performance of probabilistic forecast systems, reliability and resolution are most useful (Murphy, 1993; Winkler et al., 1996). A probabilistic forecast system is reliable if the conditional distribution of the (random) observation Y given the forecast distribution F , noted $[Y|F]$, is equal to the forecast distribution: $[Y|F] = F$. Resolution refers to the ability of the forecast system to issue forecast distributions very different from the marginal distribution of the observations. If a forecast is reliable, decisions can then be made by considering the observation will be drawn from the forecast distribution. For a reliable forecasting system, the higher is its resolution, the more useful it is useful for economical decision taking (Richardson, 2001).

In order to compare the two approaches of probabilistic forecast selection, as stated in the Introduction, two tools are used: the Continuous Ranked Probability Score which is a global measure of performance, and the rank histogram which is linked to reliability.

The Continuous Ranked Probability Score (CRPS, Matheson and Winkler 1976) is a scoring rule that quantifies the forecast performance of a probabilistic forecasts expressed as a CDF F with a scalar observation $y \in \mathcal{Y} \subseteq \mathbb{R}$

$$\ell(F, y) = \int_{x \in \mathcal{Y}} (F(x) - H(x - y))^2 dx. \quad (7)$$

The CRPS being convex in its first argument, the existence of theoretical bounds for the regret is ensured for some aggregation methods by the theory of prediction with expert advice. Since the experts are step-wise CDFs with steps at $x_{e;t}^{m_e}$, the information about the underlying forecast CDF F is incomplete, which makes some estimators of the CRPS biased, as investigated in Zamo and Naveau (2018). This last article shows that, in order to get an accurate estimate of the CRPS, one has to choose an estimator of the CRPS according to the nature of the $x_{e;t}^{m_e}$ (a sample from F or a set of quantiles from F). This recommendation was followed in this study. For a forecast CDF built as the empirical CDF of an M -random sample $x_{i=1, \dots, M}$ (such as a raw ensemble), the CRPS is estimated with

$$\widehat{CRPS}_{INT}(M) = \frac{1}{M} \sum_{i=1}^M |x_i - y| - \frac{0.5}{M^2} \sum_{i,j=1}^M |x_i - x_j| \quad (8)$$

and for a forecast CDF defined from a set of M quantiles $x_{i=1, \dots, M}$ with regularly spaced orders (such as a post-processed ensemble), the following estimator is used

$$\widehat{CRPS}_{PWM}(M) = \frac{1}{M} \sum_{i=1}^M |x_i - y| - \frac{0.5}{M(M-1)} \sum_{i,j=1}^M |x_i - x_j| \quad (9)$$

The expectation of the CRPS can be decomposed into three terms that quantify different properties of a forecasting system or the observation distribution: the

reliability (RELI), the resolution (RES) and the uncertainty (UNC) terms:

$$\mathbb{E}_{F,Y}\ell(F, Y) = RELI - RES + UNC \quad (10)$$

Let us note $\bar{\pi} = [Y]$ the marginal distribution of the observation Y and $\pi^F = [Y|F]$ the conditional distribution of the observation Y given the forecast distribution F . Then the three terms are defined as in Bröcker (2009) by

$$RELI = \mathbb{E}_{F,Y \sim \pi^F} (\ell(F, Y) - \ell(\pi^F, Y)) \quad (11)$$

$$RES = \mathbb{E}_{F,Y \sim \pi^F} (\ell(\bar{\pi}, Y) - \ell(\pi^F, Y)) \quad (12)$$

$$UNC = \mathbb{E}_{Y \sim \bar{\pi}} \ell(\bar{\pi}, Y) \quad (13)$$

The reliability term is the average difference of CRPS between forecasts F and π^F when forecasting observations distributed according to π^F . It is negatively oriented with a minimum of 0 for a perfectly reliable forecast system ($\pi^F = F \ \forall F$). The resolution term is the average difference of CRPS between forecasts $\bar{\pi}$ and π^F when forecasting observations distributed according to π^F . It is positively oriented and the more F is different from $\bar{\pi}$ the higher it is. For a reliable forecast system, it is essentially equivalent to the sharpness. The uncertainty term is linked to the dispersion of the marginal distribution of the observation and does not depend on the forecast. Hersbach (2000) gives a method to estimate these terms from an ensemble forecast. Based on this decomposition, good CRPS can be obtained with a bad reliability (high RELI) provided that the resolution is good enough (high RES). By minimizing the CRPS as a forecast selection criterion, one may thus select a forecast system that has a low CRPS but is not reliable (Wilks, 2018). The solution proposed by Wilks (2018) is to minimize the CRPS modified with a penalty proportional to the reliability term. Although this may indeed decrease the reliability term, there is no way to check if it is low enough.

Collet and Richard (2017) proposed a post-processing method that enforces the reliability of the forecast. To have a more general forecast selection method, it is proposed here to check that a forecasting system is reliable by testing the hypothesis that its rank histogram is flat. This procedure may be used as long as one can build a rank histogram from the forecast, which is generally the case. The rank histogram of an ensemble forecast, simultaneously introduced by Anderson (1996), Hamill and Colucci (1996) and Talagrand et al. (1997), is the histogram of the rank of the observation when it is pooled with its corresponding forecast members. For a reliable ensemble, the observation and the members must have the same statistical properties, resulting in a flat rank histogram. The type of deviation from flatness gives indications about the flaws of an ensemble. For instance, an L-shape histogram means the forecasts are consistently too high, while a J-shape histogram indicates consistently too low forecasts. A U-shape histogram reveals the forecast distribution is under-dispersed or conditionally biased. Hamill (2001) showed on synthetic data that a flat rank histogram can be obtained for an unreliable ensemble. Producing a flat rank histogram is thus a necessary but not sufficient condition for a forecasting system

to be reliable. Although the rank histogram was initially designed for ensemble forecasts, it can be used for forecasts in the form of quantiles. If the quantiles' orders are regularly spaced between 0 and 1 (excluded), then the histogram should also be flat if the forecasting system is reliable. The flatness of a rank histogram can be statistically tested thanks to the Jolliffe-Primo tests of flatness described in Jolliffe and Primo (2008) and summarized in Appendix A. These tests assess the existence of some specific deviations from flatness in a rank histogram, such as a slope, a convexity and a wave shape¹. In this study we apply the three flatness tests to each rank histogram at several locations where forecasts and observations are available, which leads to a multiple testing procedure. In order to take into account this multiple testing, the false discovery rate is controlled thanks to the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995; Benjamini and Yekutieli, 2001) with an alpha of 0.01. Bröcker (2018) showed that when the observation ranks are serially dependent, the JP tests should be adapted to take into account this temporal dependency. In this study, the lag-1 autocorrelation of the rank has a median of 0.2 over all the studied locations and lead times and is lower than 0.4 for most of the forecasting systems (raw, post-processed and aggregated alike). This correlation was judged low enough so the modified procedure proposed by Bröcker (2018) was not used.

3 Aggregation Methods

Five aggregation methods are introduced, from simple empirical ones to more sophisticated ones derived from the theory of prediction with expert advice.

3.1 Inverse CRPS Weighting

The *inverse CRPS weighting method* (INV) gives to each expert's forecast a weight inversely proportional to its average CRPS over the last W days:

$$\omega_{e;t} = \frac{(\overline{CRPS}_{e;t})^{-1}(W)}{\sum_{e=1}^E (\overline{CRPS}_{e;t})^{-1}(W)} \quad (14)$$

where $\overline{CRPS}_{e;t}(W)$ is the average CRPS of expert e during the W days before time t .

¹The p-values of the Jolliffe-Primo test for slope and convexity have been computed based on a modified version of the function `TestRankhist` in the R package `SpecsVerification` (Siegert, 2015). The function has been modified to compute also the p-value for the Jolliffe-Primo test for a wave shape, introduced in this study.

3.2 Sharpness-Calibration Paradigm

The *sharpness-calibration paradigm* (SHARP) of Gneiting et al. (2007) can be used as an aggregation method to select at each time one expert. The aggregated forecast is the forecast of the expert whose range of the central 90% interval IQ_{90} , averaged over the last W days is the lowest, among the experts whose reliability term, as computed in Hersbach (2000), is lower than a chosen threshold $Reli_{th}$ over the last W days. The aggregated forecast gives an aggregation weight of 1 to the sharpest reliable expert’s forecast and of 0 to the other experts’ forecast:

$$\omega_{e;t} = \mathbb{1} \left(e = \underset{\{e | Reli_{e;t}(W) < Reli_{th}\}}{\operatorname{argmin}} \overline{IQ}_{90_{e;t}}(W) \right), \quad (15)$$

where $Reli_{e;t}(W)$ is the reliability term of expert e over the W days before time t , and $\overline{IQ}_{90_{e;t}}(W)$ is the average range of the interval between quantiles of orders 0.95 and 0.05, forecasted by the expert e over the W days before time t . If no expert has a reliability term lower than the reliability threshold $Reli_{th}$ over the last W days, the aggregated forecast is just the expert with the lowest mean CRPS over the last W days.

The three following methods are derived from the theory of prediction with expert advice, and bounds for the regret may be computed.

3.3 Minimum CRPS

The *minimum CRPS method* (MIN) chooses the best recent expert in terms of CRPS, that is, the aggregation weight is 1 for the expert with the lowest average CRPS over the last W days, and 0 for all the other experts:

$$\omega_{e;t} = \mathbb{1}(e = e_t^*(W)), \quad (16)$$

where $e_t^*(W)$ is the index of the expert with the minimum average CRPS during the last W days. The reference class \mathcal{C} is the set of the E available experts, so that the oracle for this method is the expert with the lowest CRPS averaged over the period $\{1, \dots, T\}$, $e_T^*(T)$. This aggregation method is called “follow-the-best-expert” in Cesa-Bianchi et al. (2006), who prove that, under several assumptions on the loss, the regret of the aggregated forecast relatively to the oracle is $o(\ln(T))$.

3.4 Exponential Weighting

The *exponentially weighted average forecaster* (EWA) computes the aggregation weights as

$$\omega_{e;t} = \frac{\exp\{-\eta CRPS_{e;t}(W)\}}{\sum_{e=1}^E \exp\{-\eta CRPS_{e;t}(W)\}}, \quad (17)$$

where $\eta \in \mathbb{R}^+$ is called the learning rate and $CRPS_{e;t}(W)$ is the cumulative CRPS of expert e over the last W days. The highest the learning rate, the lowest the weight for a bad expert. For very large learning rates η , EWA is practically equivalent to the MIN method. The reference class \mathcal{C} is the set of the E available experts. If W spans the whole period before t , that is, if $W = t - 1$ days, EWA competes well with the best expert, in terms of average CRPS over the whole period $\{1, \dots, T\}$, with a bounded regret relatively to this oracle

$$\sup R_T^{\mathcal{C}} \leq \frac{\ln E}{\eta} + \frac{\eta T}{8} B^2, \quad (18)$$

where B is the upper bound of the loss function (see proof in Appendix B). The demonstration requires only the convexity of the loss ℓ in its first argument. In practice, for an unbounded loss such as the CRPS, B is the maximum of observations and expert forecasts over the whole period $t = 1, \dots, T$. Although the theoretical bounds exist only for $W = t - 1$, one can use a shorter sliding window W in order to make the aggregation weights change quickly over time. This may improve the forecast of an unstationary variable of interest.

3.5 Exponentiated Gradient

The *exponentiated gradient forecaster* (GRAD) weights the experts with

$$\omega_{e;t} = \frac{\exp\{-\eta \partial_e CRPS_t(W)\}}{\sum_{e=1}^E \exp\{-\eta \partial_e CRPS_t(W)\}}, \quad (19)$$

where $CRPS_t(W)$ is the cumulative CRPS of the GRAD aggregated forecast over the last W days, and $\partial_e CRPS_t(W) = \frac{\partial CRPS_t(W)}{\partial \omega_{e;t}}$. Using Appendix C leads to

$$\begin{aligned} \partial_e CRPS_t(W) &= \sum_{s=t-W}^{t-1} \frac{\partial CRPS}{\partial \omega_e}(\hat{y}_s, y_s) \\ &= \sum_{s=t-W}^{t-1} \left\{ \sum_{m_e=1}^{M_e} p_e^{m_e} |x_{e;s}^{m_e} - y_s| - \sum_{e'=1}^E \omega_{e';s} \sum_{m_{e'}=1}^{M_{e'}} p_{e'}^{m_{e'}} |x_{e';s}^{m_{e'}} - y_s| \right\} \\ &\quad - \sum_{s=t-W}^{t-1} \left\{ \sum_{e'=1}^E \omega_{e';s} \left(\sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} |x_{e;s}^{m_e} - x_{e';s}^{m_{e'}}| \right) \right\}, \end{aligned} \quad (20)$$

which generalizes Equation (5.13) of Baudin (2015) to the case of the aggregation of step-wise CDFs with any number of steps. The rationale for using the gradient is that, if the gradient is positive for an expert over the past W forecasts, increasing the weight would have increased the CRPS (and decreased

the performance) of the aggregated forecast. So for incoming forecasts, giving this expert a low weight should improve the forecast performance. The reverse is true for a negative gradient.

For this aggregation method, the reference class \mathcal{C} is the set of convex combinations with constant weights over the whole period $\{1, \dots, T\}$, that is, such that $\tilde{\omega}_{e;t} = \tilde{\omega}_e, \forall e \in \{1, \dots, E\}$ and $t \in \{1, \dots, T\}$. The oracle is the best, in terms of cumulative CRPS, constant convex combination of experts. This is usually a better oracle than the best expert. If $W = t - 1$ days, the following bound immediately follows from Mallet et al. (2007) or Baudin (2015)

$$\sup R_T^{\mathcal{C}} \leq \frac{\ln E}{\eta} + \frac{\eta T}{2} C^2, \quad (21)$$

where $C = \max_{t \in \{1, \dots, T\}, e \in \{1, \dots, E\}} |\frac{\partial CRPS}{\partial \omega_e}(\hat{y}_t, y_t)|$. Although the oracle for GRAD is better than for EWA, the bounds may be larger, so that EWA may actually perform better than GRAD. In practice, one has to try both methods to know which one performs best.

4 The Experts and the Observation

In this work, aggregation is applied to ensemble forecasts of the 10 m wind speed over France. Forecasting surface wind is quite difficult due to complex interactions between phenomena at different spatio-temporal scales. For instance, wind speed is influenced by large scale atmospheric structures such as cyclones and anticyclones, but also by local orography and surface friction. Local atmospheric effects such as downward drifts under convective clouds also contribute to the direction and speed of wind. National weather services need to have skillful wind speed forecasts to issue early warnings to the population and civil security services. Also, wind speed influences many economic activities, such as sailing, windpower generation, construction,... that need good forecasts for decision making.

This section presents the 28 experts used in this study and the wind speed observation used to assess the forecasts.

4.1 The Four Experts based on TIGGE

The International Grand Global Ensemble, formerly the THORPEX Interactive Grand Global Ensemble (TIGGE), was an international project aiming, among other things, to provide ensemble prediction data from leading operational forecast centers (Bougeault et al., 2010; Swinbank et al., 2016). Although

the TIGGE data set² includes 10 ensemble NWP models, only the four ensemble models available on a daily basis at Météo-France have been retained (see Table 1) as experts, so that the aggregation methods may later be used in operations at Météo-France.

The TIGGE ensembles are available on a grid size is 0.5°. Over France this amounts to a total of 267 grid points. The study period goes from the 1st January, 2011 to the 31st December, 2014 (so $T = 1461$ in the notations of Section 2). The lead times go from 6 h to 54 h depending on the ensemble, with a timestep of 6 h. The forecast are done at 1800 UTC for lead times h from 6 h to 48 h, with a time step of 6 h. This implies that for experts based on CMC and ECMWF, whose runtime is 1200 UTC, the actual lead time is $h + 6$ (see Table 1).

Each ensemble is an expert whose forecast CDF $\hat{F}_{e;t}$ is the empirical CDF of the members associated with the same weight $p_e^{me} = \frac{1}{M}$, where M is the number of members in the ensemble.

Table 1 – Ensembles from TIGGE used in this study, with some of their characteristics.

Weather service	Members	Hour of the run used (UTC)	Lead times
Canadian Meteorological Center (CMC)	21	1200	12h to 54h
European Center for Medium-Range Weather Forecasts (ECMWF)	51	1200	12h to 54h
Météo-France (MF)	35	1800	6h to 48h
US National Centers for Environmental Prediction (NCEP)	21	1800	6h to 48h

4.2 The Twenty Four Experts built with EMOS Methods

Each ensemble is post-processed with two kinds of EMOS: non-homogeneous regression (NR, Gneiting et al. 2005, Hemri et al. 2014) and quantile random forest (QRF, Meinshausen 2006, Zamo et al. 2014, Taillardat et al. 2016).

The forecast CDF produced with NR is parametric. Following Hemri et al. (2014), the square root of the forecast wind speed \hat{f}_t is supposed to follow a normal distribution truncated at 0

$$\sqrt{\hat{f}_t} \sim \mathcal{N}^0(a + b\bar{x}_t, c^2 + d^2 sd_t) \quad (22)$$

where \bar{x}_t and sd_t are the mean and standard deviation of the square-root of the associated ensemble values, forecasted at time t . The four parameters (a, b, c and d) are optimized by maximizing the log-likelihood³ over the last W_{tr} fore-

²The TIGGE data set can be retrieved from the ECMWF at <http://apps.ecmwf.int/datasets/> or from the Chinese Meteorological Administration at <http://wisportal.cma.gov.cn/wis/>

³With the function `optim` in R (R Core Team, 2015).

cast days. In order to build several NR-post-processed ensembles with different reaction time to the raw ensemble’s performance changes, five sizes of the training window W_{tr} are used, as summarized in Table 2. These parametric CDFs can not be used as such in the framework of step-wise CDF aggregation. A step-wise $\hat{y}_{e;t}$ is built from the squared quantiles of orders $\{0, \frac{1}{100}, \dots, \frac{99}{100}, 0.999\}$ of the parametric CDF produced by NR. The last order is not 1 in order to avoid infinite values.

QRF is a non parametric machine learning method that produces a set of quantiles of chosen order. A random forest is a set of regression trees built on a bootstrapped version of the training data (Zamo et al., 2016). While building each tree, the observations are split in two most homogeneous groups (in terms of variance of the observation), according to a splitting criterion over an explanatory variable. At each split (or node) only a random subset of the complete set of explanatory variables are tested, until a stopping criterion is reached (such as a maximum number of final nodes, called leaves). When a forecast is required, a step-wise CDF is produced by going down the forest with the vector of explanatory variables, computing the step-wise CDF of the observations associated to each leaf and averaging those CDFs. In practice, one requests a set of quantile orders and gets the corresponding quantiles⁴. The requested quantile orders are the same as for NR, except the last order which is 1, because QRF cannot produce infinite values. Since the forecast CDF is a step-wise function, the obtained quantiles may contain many ties, that are suppressed by linearly interpolating between the points defining the CDF’s steps, as explained in Zamo and Naveau (2018).

The post-processing and aggregation methods are trained separately for each of the eight lead times and each grid point. Since NR is a parametric method with only four parameters, it can be trained with few data, thus a sliding-window training period is used. QRF, being a non parametric method, requires more training data to define the many splits in each tree. It is thus trained with 4-fold cross-validation: each year is successively used as a test sample, the remaining three being used as the training sample.

From each of the four ensembles seven experts are built: the raw ensemble, the QRF-post-processed ensemble and five NR-post-processed versions of the ensemble (each with a different size of the training window, Table 2). Altogether, 28 experts are aggregated.

4.3 The Observation

The observation is the 10 m average wind speed analysis built in Zamo et al. (2016), at 267 grid points over France. Those grid points are the same as the ensemble datasets presented in Section 4.1.

⁴In the R packages `quantregForest` or `ranger`.

Table 2 – Summary description of the EMOS methods used to post-process each ensemble.

Quantile Regression Forest (QRF)	
Forecast distribution	Non-parametric (set of quantiles).
Explanatory variables	Control member, ensemble mean, ensemble 0.1 and 0.9 quantiles, month.
Training method	4-fold cross-validation (3 training years, 1 test year).
Orders of the forecast quantiles	$0, \frac{1}{100}, \dots, \frac{99}{100}, 1$.
Non-homogeneous Regression (NR)	
Forecast distribution	Parametric (truncated normal distribution for the square-root of wind speed).
Explanatory variables	Mean and standard deviation of the raw ensemble.
Training method	Likelihood maximization on a sliding window over the W_{tr} previous days. Five windows are used: $W_{tr} = 7, 30, 90, 365, t - 1$ days
Orders of the forecast quantiles	$0, \frac{1}{100}, \dots, \frac{99}{100}, 0.999$.

5 Results

The five aggregation methods presented in Section 3 have been investigated, with different values of the tuning parameters W and η . The values of the tuning parameters tested in the study are listed in Table 3. When two parameters are listed, all combinations have been tried.

Table 3 – Sets of tried values for the parameters of the aggregation methods. When two parameters are given, all the possible combinations of values have been tried.

Aggregation method	Parameters' values
Minimum CRPS or Inverse CRPS	$W = 7, 15, 30, 90, 365, t - 1$ days
Sharpness-calibration	$W = 7, 15, 30, 90, 365, t - 1$ days, $Reli_{th} = 0.1$ m/s
Exponentiated weighting or Exponentiated gradient weighting	$W = 7, 15, 30, 90, 365, t - 1$ days, $\eta = 10^{-1.5}, 10^{-1}, 10^{-0.5}, 1, 10^{0.5}, 10^{1.5}, 10^2$

Hereafter the best expert and the best aggregation method is selected following the two approaches mentioned above: the minimization of the average CRPS or the maximization of the proportion of grid points with a rank histogram deemed flat by the three flatness tests. The best forecasting system in terms of

minimum CRPS is called the most skillful. The forecasting system which gets the maximum proportion of grid points with a flat rank histogram is called the most reliable.

5.1 CRPS and Reliability

From Table 4, the most skillful expert is the QRF-post-processed ECMWF ensemble (it is an oracle). The time series of the regret of each most skillful or reliable aggregation method of each type relatively to the QRF-post-processed ECMWF ensemble is drawn in Figure 2, for lead time 24 h. Whereas the most skillful setting of the SHARP aggregation method gets a consistently higher CRPS than the most skillful expert, the other aggregation methods manage to outperform the latter at least for some part of the four years. The most skillful settings of EWA and GRAD even get a negative regret relatively to the most skillful expert. The regret exhibits a trend and a diurnal cycle with the lead times (not shown), and so does the averaged CRPS (see Table 4): the forecast performance decreases with increasing lead times, particularly during the late afternoon when the wind strengthens. The main point is that, in terms of CRPS, post-processing improves performance, and aggregation further improves performance. According to the minimization of the CRPS, the chosen forecast method would be the most skillful GRAD setting, that is $\log_{10}(\eta) = -1$ and $W = t - 1$ days.

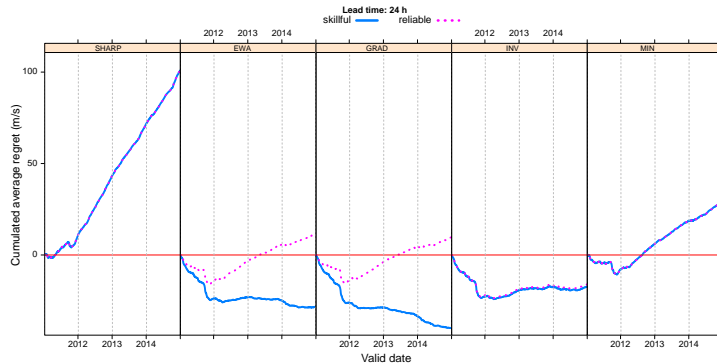


Figure 2 – Time series of the cumulative spatially-averaged regret, at lead time 24 h, for each aggregation method, relative to the most skillful expert (QRF-post-processed ECMWF ensemble). At each valid date, the regret relatively to the QRF-post-processed ECMWF ensemble is averaged over the 267 grid-points. For each aggregation method, two settings are used to compute the regret: the most skillful one (blue continuous line) and the most reliable one (pink dashed line).

If one uses the flatness of rank histograms as the selection criterion, the best raw ensemble is CMC, the best expert is QRF-post-processed MF and the best

Table 4 – CRPS averaged over the four years and the 267 grid points for several forecasting systems. The average CRPS of the most skillful raw ensemble (ECMWF) is indicated in the first line. Then, for each model selection approach, the best expert is indicated followed by the best setting of each aggregation method. For each selection approach, the forecasting system with the lowest CRPS is in bold.

Method	Parameters		Lead time (h)								
	$\log_{10}(\eta)$	W	all	6	12	18	24	30	36	42	48
RAW											
ECMWF			0.76	0.79	0.79	0.73	0.73	0.78	0.78	0.73	0.74
Selection: most skillful forecasting system.											
expert											
(QRF			0.49	0.47	0.46	0.48	0.50	0.49	0.49	0.52	0.53
ECMWF)											
SHARP		1095	0.55	0.52	0.52	0.55	0.56	0.55	0.55	0.59	0.60
GRAD	-1	t-1	0.47	0.44	0.44	0.46	0.47	0.47	0.47	0.51	0.51
EWA	-1	365	0.48	0.44	0.44	0.47	0.48	0.47	0.48	0.52	0.52
INV		7	0.49	0.46	0.46	0.47	0.48	0.49	0.49	0.52	0.52
MIN		365	0.51	0.47	0.47	0.51	0.52	0.50	0.51	0.56	0.56
Selection: most reliable forecasting system.											
expert											
(QRF			0.50	0.46	0.47	0.51	0.50	0.49	0.50	0.55	0.53
MF)											
SHARP		1095	0.55	0.52	0.52	0.55	0.56	0.55	0.55	0.59	0.60
GRAD	0.5	t-1	0.50	0.46	0.46	0.50	0.50	0.49	0.50	0.54	0.54
EWA	0.5	30	0.50	0.46	0.46	0.50	0.50	0.50	0.50	0.55	0.54
INV		7	0.49	0.46	0.46	0.47	0.48	0.49	0.49	0.52	0.52
MIN		365	0.51	0.47	0.47	0.51	0.52	0.50	0.51	0.56	0.56

aggregated forecast is EWA with $\log_{10}(\eta) = 0.5$ and $W = 30$ days (see Table 5). Raw CMC is not a very reliable ensemble, as shown in Figure 3 (a) for lead time $h = 6$ h: the ensemble is consistently biased with too strong forecast wind speeds in the north-west of France and too weak forecast wind speeds over the Alps and the Pyrenees. Elsewhere, although the ensemble is much less biased, the rank histogram is not deemed flat due to an obvious U-shape. The rank histograms at the other lead times and for the other raw ensembles exhibit similar features (not shown). For the most reliable expert and aggregated forecasts, the rank histograms are computed with the nine forecast deciles. As illustrated in Figure 3 (b), the QRF-post-processed version of the MF ensemble yields a higher number of flat rank histograms than the raw CMC ensemble. Finally, Figure 3 (c) shows that, the JP tests do not reject the flatness hypothesis at many more grid-points for the most reliable EWA forecast. Table 5 confirms quantitatively that the most reliable EWA outperforms the most reliable expert in terms of flatness of the rank histogram, at each lead time. The most reliable

Table 5 – Proportion of grid points where the three flatness tests do not reject the hypothesis of a flat rank histogram for several forecasting systems. For each model selection approach, the best expert is indicated followed by the best setting of each aggregation method. The forecasting system with the highest proportion of rank histograms is in bold, for each selection approach.

Method	Parameters		Lead time (h)								
	$\log_{10}(\eta)$	W	all	6	12	18	24	30	36	42	48
Selection: most skillful forecasting system.											
expert											
(QRF ECMWF)			0.60	0.26	0.17	0.83	0.89	0.41	0.43	0.95	0.86
SHARP		1095	0.39	0.28	0.22	0.39	0.36	0.53	0.45	0.45	0.41
GRAD	-1	t-1	0.05	0.06	0.06	0.03	0.04	0.06	0.05	0.03	0.04
EWA	-1	365	0.04	0.04	0.03	0.03	0.04	0.06	0.06	0.05	0.04
INV		7	0.01	0.03	0.02	0.00	0.01	0.02	0.01	0.00	0.01
MIN		365	0.82	0.83	0.79	0.84	0.69	0.88	0.87	0.92	0.75
Selection: most reliable forecasting system.											
expert											
(QRF MF)			0.92	0.76	0.89	0.99	0.99	0.87	0.87	1.00	0.99
SHARP		1095	0.39	0.28	0.22	0.39	0.36	0.53	0.45	0.45	0.41
GRAD	0.5	t-1	0.67	0.46	0.43	0.91	0.79	0.46	0.51	0.98	0.79
EWA	0.5	30	0.97	0.96	0.90	0.98	0.99	0.97	0.97	0.98	0.98
INV		7	0.01	0.03	0.02	0.00	0.01	0.02	0.01	0.00	0.01
MIN		365	0.82	0.83	0.79	0.84	0.69	0.88	0.87	0.92	0.75

setting of the other aggregation methods produce fewer flat rank histograms than the most reliable expert, for all lead times, as shown in Table 5. In other words, post-processing improves reliability over raw ensemble, and aggregation may further improve reliability over the most reliable expert if the right setting is chosen.

The two selection approaches choose a different “best” forecast. The GRAD aggregation method with $\log_{10}(\eta) = -1$ and $W = t - 1$ days is the most skillful forecast, whereas EWA with $\log_{10}(\eta) = 0.5$ and $W = 30$ days is the most reliable forecast. In terms of averaged CRPS over all lead times, the most skillful forecast is about 6 % more performant than the most reliable forecast, as computed from Table 4 (0.47 m/s versus 0.50 m/s). Table 5 shows that the most skillful forecast passes the three flatness tests for only 5% of grid points, whereas the most reliable forecast produces about 97% of flat rank histograms. In other words, selecting the most reliable forecast as best forecast, instead of the most skillful, leads to a slightly worse CRPS but increases dramatically the number of grid points with a flat rank histogram. This discrepancy between the two criteria for choosing the best forecasts is discussed more deeply in Section 6.

5.2 Temporal Variation of the Weights

The most reliable EWA aggregation is able to quickly redistribute the aggregation weights between the experts, as illustrated in Figure 4. For instance, during the middle of 2011, nearly all the aggregation weight shifts from the QRF-post-processed MF ensemble to the QRF-post-processed ECMWF ensemble, in a few days. The aggregation weights can also remain stable for long periods of time, such as around mid-2014, when the QRF-post-processed MF ensemble keeps a high weight for about two months. Moreover, although the raw ensembles do not perform well on their own, the EWA method may find periods where the raw ensembles can significantly contribute to the aggregated forecasts, such as in late 2012 for the raw CMC ensemble. Last, the time series of the aggregation weights is very different from one lead time to another (not shown). These features prove that this aggregation method is very adaptive. This may be very useful for operations when an ensemble undergoes important changes: the aggregation method will quickly detect a modification in performances and adjust its weighting accordingly.

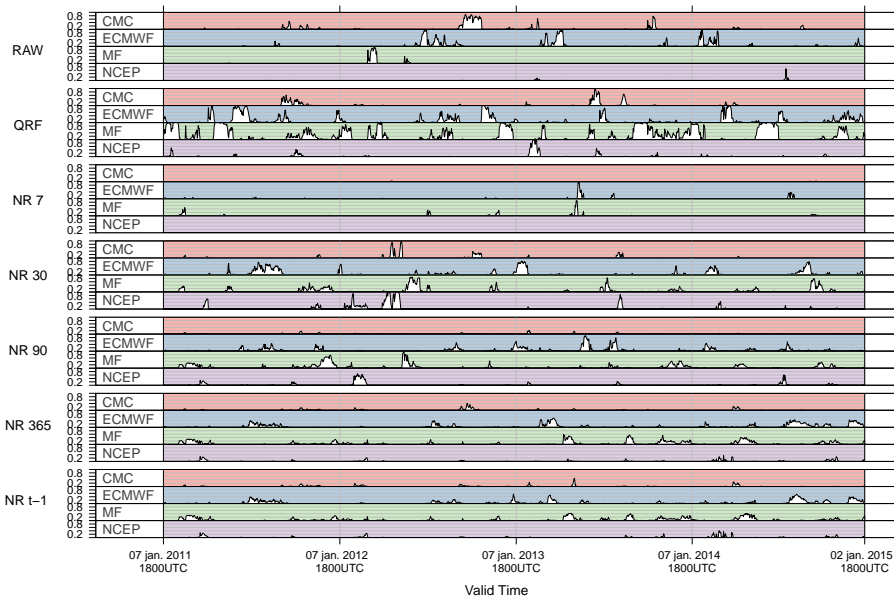


Figure 4 – Evolution of the aggregation weights with the valid date, for lead time 42 h. The aggregation method is the EWA forecaster with $\log_{10}(\eta) = 0.5$ and $W = 30$ days.

5.3 Aggregation of individual sorted experts

Baudin (2015) and Thorey (2017) defines experts as the sorted values of the pooled raw ensembles, and aggregate step-wise CDFs with one step. To compare our approach to their’s, we aggregated the sorted values of the raw ensembles. This amounts to 128 experts. We also aggregated the sorted values of the raw and post-processed ensembles. This amounts to 2,552 experts (the 128 values forecasted by the raw ensembles plus 101 quantiles for each of the 24 post-processed forecasts). A comprehensive comparison with the approach in Baudin (2015) being out of the scope of this article, only the most reliable aggregation method (EWA, with $\log_{10}(\eta) = 0.5$ and $W = 30$ days) has been used.

The resulting CRPSs are shown in Table 6. The forecast skill is improved compared to the raw ensembles (compare with Table 4). Even if no post-processed ensemble is used, the CRPS is improved by the aggregation. But adding more experts (among which are post-processed ensembles) further increases the performance improvement. However, the CRPS stays higher than the CRPS of the EWA aggregation of multi-step CDFs (0.52 m/s versus 0.50 m/s with our most reliable EWA). Although the differences in CRPS is low, the difference between the two approaches proves much more important in terms of proportion of flat rank histograms. Whereas our most reliable EWA gets about 97% of flat rank histograms (see Table 5), the aggregation with EWA of one-step CDFs gets no flat rank histogram whatever lead time is considered. All the p-values of the Jolliffe-Primo tests are lower than $2 \cdot 10^{-25}$ (with a lag-1 autocorrelation of the rank around 0.12). Further investigations would be required to check if this holds for other settings of EWA and for other aggregation methods. A possible explanation of this lack of reliability would be the fact that, for 1-member ensembles, the CRPS reduces to the mean absolute error (MAE). Since the MAE is minimized by the conditional median of the observation given the forecast, each expert is weighted according to its ability to predict well the conditional median, which is not what is looked for.

Table 6 – CRPS of the aggregated forecast built from experts considered as one-step CDFs. The experts are either the sorted members of the raw ensembles, or the sorted values of the raw and post-processed ensembles. The aggregation method is EWA with $\log_{10}(\eta) = 0.5$ and $W = 30$ days.

Aggregated experts	Lead time (h)								
	all	6	12	18	24	30	36	42	48
Sorted raw members ($E = 128$)	0.58	0.57	0.57	0.57	0.57	0.58	0.58	0.61	0.60
Sorted raw and post-processed members ($E = 2,552$)	0.52	0.48	0.49	0.50	0.51	0.52	0.53	0.56	0.56

6 Discussion about Probabilistic Forecast Selection

The discrepancy between the choice of the best forecast according to the CRPS and according to the flatness of rank histograms is now more fully investigated and discussed.

Although the CRPS is a natural measure of performance for forecast CDF, minimizing it does not ensure to get the highest number of grid-points with a flat histogram, as stated earlier and as confirmed in Figure 5 (left). In this figure, for each lead time, EWA and GRAD reach a minimum average CRPS for very low proportions of flat rank histograms. Actually, the point corresponding to this optimal CRPS is indicated by the graph of INV in Figure 5 (left). The graph for INV seems to reduce to a point because the associated CRPS and proportion of flat rank histograms barely vary. For post-processed ensembles, the same behavior can be observed (see Appendix D): the most skillful expert may not be the most reliable one.

This discrepancy can be explained with the decomposition of the average CRPS as a sum of a reliability term (that must be minimized) minus a resolution term (that must be maximized) and an uncertainty term (that does not depend on the forecasting system). The reliability and the resolution terms in the average CRPS after post-processing or aggregation have very different ranges of variation from one setting to another, as deduced from Figure 5 for the aggregated forecasts. For instance, at a fixed lead time, the reliability term of the MIN methods varies by less than 0.01 m/s while the CRPS varies by 0.05 m/s. Since the uncertainty term depends only on the observation and is fixed for a fixed lead-time, the CRPS varies mainly because of the variations in the resolution term. In other words, in this case, the average CRPS is mostly a measure of resolution and its minimization as a selection criterion leads to maximize the resolution. But the distributions may not be reliable since a small change in the reliability term will not change a lot the average CRPS while being compatible with unreliable flat rank histograms according to the JP-tests (see right side of Figure 5). This comment is also true for the other aggregation methods. It also holds for post-processed ensembles, whose reliability term may change by about 0.01 m/s whereas the average CRPS varies by 0.2 m/s with very different proportions of reliable grid points (see Appendix D).

Another part of the origin of the discrepancy between the two selection approaches is that the CRPS quantifies the forecast performance over the whole distribution, while the JP tests assess only forecast performance for the tested shapes of the rank histogram. Firstly, the rank histogram takes only into account the ranking of the observation and the members (or quantiles) while the reliability term of the CRPS takes into account the distance between the observation and the members (or quantiles). Secondly, limiting the flatness criterion to the absence of slope, convexity and wave shapes may lead to miss other important

deviations to flatness in the rank histogram. One could use more flatness tests to add constraints on the forecasting system selected on a reliability criterion. It would be interesting investigate how it changes the solution.

Finally, the sampling noise may add further differences between the solutions selected by the two approaches. The CRPS and the flatness tests may not react in the same way to this sampling. The JP-test being hypothesis testing it has the same limit as every tests. Rejection of the flatness hypothesis may be due to other features than a lack of reliability (Wasserstein et al., 2019).

In conclusion, the selection of the best forecast should not be made only by minimizing the CRPS, but also by taking care of the actual reliability of the forecasts, as tested with the JP tests. Both performance criteria should be used. Being only focused on the CRPS may lead to choosing a forecasting system that is not optimal in terms of reliability as shown here and earlier studies (Collet and Richard, 2017; Wilks, 2018). But relying only on the JP flatness tests may also be misleading. Indeed, always forecasting the climatology leads to a very good reliability but a very low resolution. Consequently, a forecaster may be tempted to forecast the climatology in order to get a good forecast, instead of issuing a forecast he might consider more likely but too different from the climatology, and too risky to issue. Choosing a forecast based on the JP tests only may lead to such hedging strategies. In this study, hedging is avoided by using the CRPS to tune the experts and the aggregation. But the reliability of the chosen forecast is ensured by using the JP tests.

7 Conclusion and Perspectives

The first goal of the present study was to adapt the theory of prediction with expert advice to the case of experts issuing probabilistic forecasts as step-wise CDFs with any number of steps. Contrary to the work of Baudin (2015) who aggregated unidentifiable experts built by pooling and sorting members of several ensembles, each expert used in the present work is identifiable over time as required by the theoretical framework of prediction with expert advice: the aggregation weights for the members or quantiles of the same expert are constrained to be equal. Some formulae of Baudin (2015) valid for step-wise CDFs with one step have been generalized to the case of step-wise CDFs with any number of steps.

Several aggregation methods to combine step-wise forecast CDFs have been presented and compared in terms of reliability and CRPS. The reliability has been assessed by using the Jolliffe-Primo tests, which detect the presence in the rank histogram of typical deviations from flatness. The systematic use of the Jolliffe-Primo flatness test highlights that the minimization of the CRPS as the main criterion to calibrate or aggregate may not produce the maximum number of flat rank histograms. It is also shown that choosing the best forecast

by maximizing the proportion of rank histograms ensures reliable forecasts, without significantly increasing the CRPS.

On a real wind speed data set, the best aggregation method, in terms of proportion of flat rank histograms, is the exponentially weighted average forecaster, with a learning rate $\eta = 10^{0.5}$ and an aggregation window $W = 30$ days. This aggregated forecast has a similar CRPS as the most skillful expert in terms of CRPS, and produces many more flat rank histograms than the most reliable expert. The method can produce weights with very different temporal patterns: rapidly evolving weighting of the experts, long period of constant weighting, short period with large weights for the raw ensembles. With the use of experts fitted over sliding windows of different size, this flexibility may help to solve a recurrent problem in post-processing: important changes in the NWP models that may make the post-processing equation inadequate for the new version of the NWP model. Although EWA is the selected aggregation method in this study, this must not be taken as a result valid for other observations and/or geographical domains. To the best of our knowledge, no theoretical results allow us to tell among the many available aggregation methods which one will be the best on a given dataset.

As for the perspectives, it is planned to study the same and other aggregation methods by pooling data in blocks of nearby grid-points. This may improve the fit by enlarging the training sample or, at the very least, speed up operations on finer grids with thousands of points, as was demonstrated for deterministic forecasts (Zamo et al., 2016). Since post-processing of other meteorological parameters, such as temperature and rainfall, has been already tested internally at Météo-France, aggregation methods will be tried on these variables too. A more comprehensive study of the discrepancy between the CRPS and the proportion of flat rank histograms as a performance criterion and its implication on post-processing and aggregation constitutes a more theoretical perspective. A further line of research would be to expand on the proposed approach to post-processing and aggregation. Indeed, the retained criterion in this study (maximizing the number of grid points with a flat rank histogram) does not allow us to choose a different forecast for each grid point, which may be desirable to further improve forecast performance.

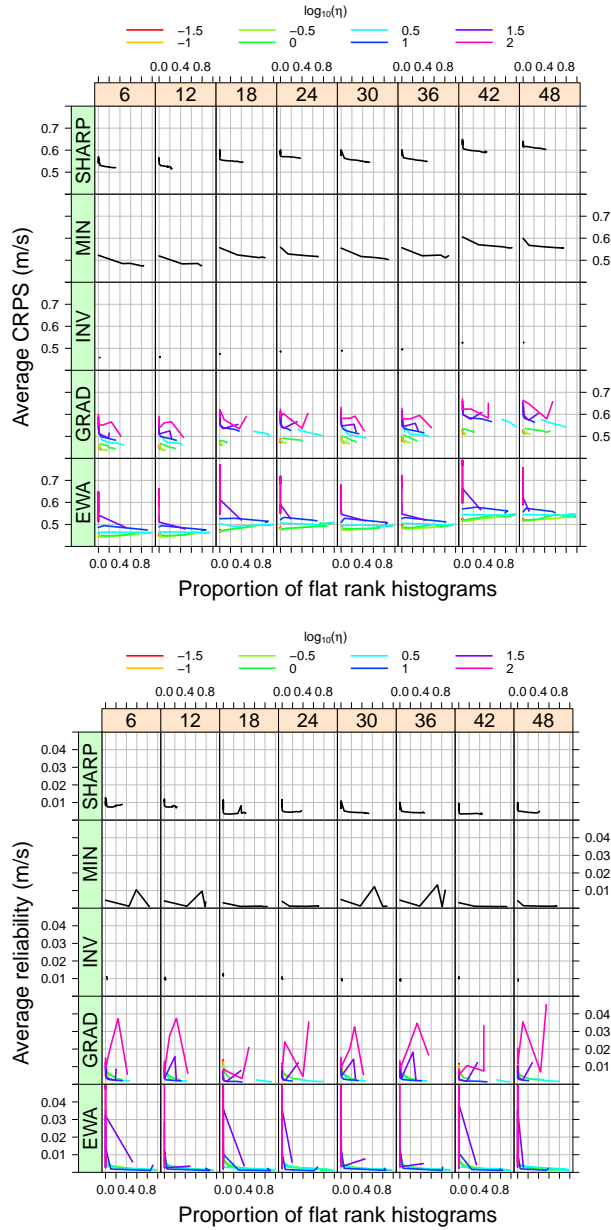


Figure 5 – CRPS averaged over space and time (left) and reliability term (right), versus the proportion of flat rank histograms. Each row corresponds to one aggregation method. Each column corresponds to one lead time. Inside each individual panel, the line is obtained by varying the aggregation window W . For EWA and GRAD, several lines are drawn, each associated with a different value of η (as indicated in the upper color legend). For INV, the variations are so small that it appears almost as a point.

Appendix

A Statistical Tests of Flatness of a Rank Histogram

Consider the vector δ of normalized deviation from flatness in each rank of some rank histogram at hand,

$$\delta = \left(\frac{n_1 - n_0}{\sqrt{n_0}}, \dots, \frac{n_k - n_0}{\sqrt{n_0}} \right),$$

where k is the number of possible ranks, n_i is the count of rank i and $n_0 = \frac{\sum_{i=1}^k n_i}{k}$ is the theoretical count in each rank for a flat histogram.

Under the null hypothesis H_0 that the rank histogram is compatible with a flat histogram up to sampling noise, the squared norm $\|\delta\|^2$ follows a χ^2 distribution with $k - 1$ degrees of liberty. The flatness of the rank histogram can be tested with a chi-square test.

The chi-square test statistic $\|\delta\|^2$ is insensitive to the shape of the deviations to a flat histogram, as shown in Figure 6. To build this figure, as in Elmore (2005), 60 integer values from 1 to 16 have been drawn from a uniform distribution. Four histograms are shown: the histogram computed with the raw sample (top left), with the same counts sorted in ascending order (top right), with the counts reassigned to have a peak-shaped histogram (bottom left), and with the counts reassigned in a wave shape (bottom right). The p-value of the chi-square test and three other flatness tests presented below is reproduced under the histograms. Although the counts of each rank are reorganized, the p-value of the chi-square test of the four histograms is the same, because reordering the counts is equivalent to reordering the components of δ , which does not change its norm. Because of this, in our study, the flatness of each rank histogram is assessed with the decomposition of the chi-square test statistic, as detailed in Jolliffe and Primo (2008). Under the null hypothesis H_0 , any projection of δ onto an orthonormal basis of \mathbb{R}^k has $k - 1$ components whose squares are asymptotically independent χ^2 random variables, each with 1 degree of freedom. If the basis vectors are chosen to describe a sloped histogram, a convex histogram, or any other shape of interest, the existence of the shape in the rank histogram can be tested. The existence of a shape is not rejected if the projection of δ onto the corresponding basis vector has a component statistically different from 0.

Jolliffe and Primo (2008) give formulae to compute the basis vectors for deviations from flatness commonly encountered on real data. As an example, if $k = 2p + 1$, the basis vector for the slope (resp. convexity) test is proportional to $(-p, -p+1, \dots, -p+(k-1))$ (resp. $(p^2 - p\frac{(p+1)}{3}, (p-1)^2 - p\frac{(p+1)}{3}, \dots, -p\frac{(p+1)}{3}, 1 -$




Figure 6 – Illustration of the Jolliffe-Primo flatness tests of a rank histogram. The p-value of four flatness tests (chi-square, and three Jolliffe-Primo tests) are reproduced.

$p^{\frac{(p+1)}{3}}, \dots, p^2 - p^{\frac{(p+1)}{3}})$). In our study, three tests are used, the slope and convexity tests, and the “wave” test not described in Jolliffe and Primo (2008). This last test assesses the presence of a deviation from flatness in the shape of a tilde, that was frequently observed in the literature (Scheuerer et al., 2015; Baran and Lerch, 2016; Taillardat et al., 2016) and in internal studies at Météo-France. The corresponding basis vector is built thanks to the Gram-Schmidt process as follows: the vector $(0, \sin(\frac{2\pi}{k-1}), \sin(2\pi\frac{2}{k-1}), \dots, \sin(2\pi\frac{k-2}{k-1}), 0)$ is made orthogonal to the slope basis vector, and the resulting vector is normalized to get the basis vector for testing the presence of a wave shape. In Figure 6, the p-values for the test of existence of a slope, a convexity or a wave are in agreement with the shape of the histograms. For instance, the low p-value of the slope test (top right) rejects flatness against slope as expected.

B Proof of the Bounds for the Regret of the Exponentially Weighted Average Forecaster

The proof closely follows the proof of theorem 2.2 in Cesa-Bianchi et al. (2006).

Let $\ell : \hat{\mathcal{Y}} \times \mathcal{Y} \rightarrow [a; b]$ be a real-valued, bounded loss function. ℓ is supposed convex in its first argument.

The EWA weights at time t are computed as

$$\omega_{e;t}^{EWA} = \frac{\exp\{-\eta L_{e;t}\}}{\sum_{e=1}^E \exp\{-\eta L_{e;t}\}},$$

with $L_{e;t} = \sum_{s=1}^{t-1} \ell(\hat{y}_{e;s}, y_s)$ the cumulative loss of expert e at time t , and with the convention $L_{e;1} = 0$ so that $\omega_{e,1}^{EWA} = \frac{1}{E} \quad \forall e$.

Let us define $W_t = \sum_{e=1}^E \exp\{-\eta L_{e;t}\} \forall t \geq 1$ and $W_0 = E$. At all times $t = 1, \dots, T$, and using the convention that a sum over 0 elements is 0 (for $t = 1$, such that $\omega_{e,0}^{EWA} = \frac{1}{E} \quad \forall e$),

$$\begin{aligned} \ln \frac{W_t}{W_{t-1}} &= \ln \frac{\sum_{e=1}^E \exp\{-\eta \ell(\hat{y}_{e;t}, y_t)\} \exp\{-\eta L_{e;t-1}\}}{\sum_{e'=1}^E \exp\{-\eta L_{e';t-1}\}} \\ &= \ln \frac{\sum_{e=1}^E \omega_{e;t-1}^{EWA} \exp\{-\eta \ell(\hat{y}_{e;t}, y_t)\}}{\sum_{e'=1}^E \omega_{e';t-1}^{EWA}}. \end{aligned} \quad (23)$$

The proof now needs Hoeffding's inequality (Hoeffding, 1963). Let $a, b \in \mathbb{R}$ with $a < b$. Let Z be a bounded random variable with values in $[a; b]$, then, $\forall s \in \mathbb{R}$, Hoeffding's inequality states that

$$\ln \mathbb{E} [e^{sZ}] \leq s\mathbb{E}[Z] + \frac{s^2}{8}(b-a)^2.$$

Using Equation (23) and Hoeffding's inequality for the random variable Z taking the values $\ell(\hat{y}_{e;t}, y_t)$ with discrete probability $\omega_{e;t-1}^{EWA}$, taking $s = -\eta$ and summing over $t = 1, \dots, T$ leads to

$$\begin{aligned} \ln \frac{W_T}{W_0} &\leq -\eta \sum_{t=1}^T \sum_{e=1}^E \omega_{e;t}^{EWA} \ell(\hat{y}_{e;t}, y_t) + \frac{\eta^2}{8}(b-a)^2 T \\ &\leq -\eta \sum_{t=1}^T \ell \left(\sum_{e=1}^E \omega_{e;t}^{EWA} \hat{y}_{e;t}, y_t \right) + \frac{\eta^2}{8}(b-a)^2 T \\ &= -\eta \sum_{t=1}^T \ell(\hat{y}_t, y_t) + \frac{\eta^2}{8}(b-a)^2 T, \end{aligned}$$

after using the convexity of the loss function ℓ in its first argument, and the definition of the EWA forecast.

Noting that the following relationship also holds

$$\begin{aligned} \ln \frac{W_T}{W_0} &= \ln \left(\sum_{e=1}^E \exp\{-\eta L_{e;T}\} \right) - \ln E \\ &\geq \ln \left(\max_{e=1, \dots, E} \exp\{-\eta L_{e;T}\} \right) - \ln E \\ &= -\eta \min_{e=1, \dots, E} L_{e;T} - \ln E, \end{aligned}$$

and combining it with the previous relationship leads to

$$-\eta \min_{e=1, \dots, E} L_{e;T} - \ln E \leq -\eta \sum_{t=1}^T \ell(\hat{y}_t, y_t) + \frac{\eta^2}{8}(b-a)^2 T.$$

Finally, dividing by $-\eta$ results in the following bound of the regret of the aggregated forecast relatively to the best expert

$$\sum_{t=1}^T \ell(\hat{y}_t, y_t) - \min_{e=1, \dots, E} L_{e;T} \leq \frac{\ln E}{\eta} + \frac{\eta}{8}(b-a)^2 T. \quad (24)$$

Noting that this bound for the regret holds for any bounded loss function ℓ convex in its first argument, which is a property of the CRPS, concludes the demonstration.

C Formula for the Gradient of the CRPS

Baudin (2015) considers the aggregation of step-wise CDFs with one single step ($M_e = 1 \quad \forall e \in \{1, \dots, E\}$). We generalize equations (5.10) and (5.13) of Baudin (2015) for, respectively, the CRPS and gradient thereof, to an aggregation of step-wise CDFs with any number of steps.

Dropping the time index t in the notations, the aggregated CDF at time t is

$$\hat{y}(x) = \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} H_{x_e^{m_e}}(x) \right],$$

with the notation $H_a(x) = H(x - a)$.

Therefore, the CRPS of the aggregated CDF at time t is

$$\begin{aligned} CRPS(\hat{y}, y) &= \int_{\mathbb{R}} \left\{ H_y(x) - \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} H_{x_e^{m_e}}(x) \right] \right\}^2 dx \\ &= \int_{\gamma}^{\Gamma} \left\{ H_y(x) - \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} H_{x_e^{m_e}}(x) \right] \right\}^2 dx, \end{aligned}$$

where $\gamma = \min(y, x_1^1, \dots, x_E^{M_E})$ and $\Gamma = \max(y, x_1^1, \dots, x_E^{M_E})$.

By developing the square inside the integral,

$$\begin{aligned}
CRPS(\hat{y}, y) &= \int_{\gamma}^{\Gamma} H_y(x) dx \\
&\quad - 2 \int_{\gamma}^{\Gamma} \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} H_{x_e^{m_e}}(x) H_y(x) \right] dx \\
&\quad + \int_{\gamma}^{\Gamma} \left\{ \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} H_{x_e^{m_e}}(x) \right] \right\} \left\{ \sum_{e'=1}^{E'} \omega_{e'} \left[\sum_{m_{e'}=1}^{M_{e'}} p_{e'}^{m_{e'}} H_{x_{e'}^{m_{e'}}}(x) \right] \right\} dx.
\end{aligned}$$

Noting that $H_a(x)H_b(x) = H_{\max(a,b)}(x)$, and $\int_{\gamma}^{\Gamma} H_a(x) dx = \Gamma - a \quad \forall a \in [\gamma; \Gamma]$, then

$$\begin{aligned}
CRPS(\hat{y}, y) &= \Gamma - y \\
&\quad - 2 \sum_{e=1}^E \omega_e \left\{ \sum_{m_e=1}^{M_e} p_e^{m_e} [\Gamma - \max(x_e^{m_e}, y)] \right\} \\
&\quad + \sum_{e,e'=1}^E \omega_e \omega_{e'} \left\{ \sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} [\Gamma - \max(x_e^{m_e}, x_{e'}^{m_{e'}})] \right\} \\
&= -y \\
&\quad + 2 \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} \max(x_e^{m_e}, y) \right] \\
&\quad - \sum_{e,e'=1}^E \omega_e \omega_{e'} \left\{ \sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} \max(x_e^{m_e}, x_{e'}^{m_{e'}}) \right\},
\end{aligned}$$

because $\sum_{e=1}^E \omega_e = 1$, and $\sum_{m_e=1}^{M_e} p_e^{m_e} = 1 \quad \forall e \in \{1, \dots, E\}$.

Since $\max(a, b) = \frac{1}{2}(a + b + |a - b|)$,

$$\begin{aligned}
CRPS(\hat{y}, y) &= -y \\
&+ \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} (x_e^{m_e} + y + |x_e^{m_e} - y|) \right] \\
&- \frac{1}{2} \sum_{e, e'=1}^E \omega_e \omega_{e'} \left\{ \sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} (x_e^{m_e} + x_{e'}^{m_{e'}} + |x_e^{m_e} - x_{e'}^{m_{e'}}|) \right\} \\
&= \sum_{e=1}^E \omega_e \left[\sum_{m_e=1}^{M_e} p_e^{m_e} (x_e^{m_e} + |x_e^{m_e} - y|) \right] \\
&- \frac{1}{2} \sum_{e, e'=1}^E \omega_e \omega_{e'} \left\{ \sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} (x_e^{m_e} + x_{e'}^{m_{e'}} + |x_e^{m_e} - x_{e'}^{m_{e'}}|) \right\}.
\end{aligned} \tag{25}$$

The derivation with respect to ω_e results in

$$\begin{aligned}
\frac{\partial CRPS}{\partial \omega_e}(\hat{y}, y) &= \sum_{m_e=1}^{M_e} p_e^{m_e} (x_e^{m_e} + |x_e^{m_e} - y|) \\
&- \sum_{e'=1}^E \omega_{e'} \left\{ \sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} (x_e^{m_e} + x_{e'}^{m_{e'}} + |x_e^{m_e} - x_{e'}^{m_{e'}}|) \right\}.
\end{aligned}$$

Finally, recalling that $\sum_{e=1}^E \omega_e = 1$, and $\sum_{m_e=1}^{M_e} p_e^{m_e} = 1 \quad \forall e \in \{1, \dots, E\}$

$$\begin{aligned}
\frac{\partial CRPS}{\partial \omega_e}(\hat{y}, y) &= \sum_{m_e=1}^{M_e} p_e^{m_e} |x_e^{m_e} - y| - \sum_{e'=1}^E \omega_{e'} \sum_{m_{e'}=1}^{M_{e'}} p_{e'}^{m_{e'}} x_{e'}^{m_{e'}} \\
&- \sum_{e'=1}^E \omega_{e'} \left\{ \sum_{m_e=1}^{M_e} \sum_{m_{e'}=1}^{M_{e'}} p_e^{m_e} p_{e'}^{m_{e'}} |x_e^{m_e} - x_{e'}^{m_{e'}}| \right\}.
\end{aligned} \tag{26}$$

Formulae (25) and (26) generalize equations (5.10) and (5.13), respectively, of Baudin (2015).

D Relationship between the experts' CRPS and the reliability term with the JP tests

The expert with the minimum CRPS may not be the most reliable, according to the proportion of flat rank histograms in view of the JP tests. For instance

in Figure 7, at lead time 36 h, the QRF-post-processed ECMWF ensemble has the lowest CRPS among the experts, but exhibits less than 40% of flat histograms despite the low significance threshold chosen for the JP tests that allows important departure from flatness. At the same lead time, the NR-post-processed version of the CMC ensemble (with $W = 90$ days) has nearly 80% of flat histograms for a slightly higher CRPS.

As for the aggregation methods, the CRPS of the post-processed ensembles is mainly driven by the resolution term. Whereas the reliability term varies by less than 0.01 m/s, the average CRPS may change by up to 0.1 m/s from one to the other. For post-processed ensembles, the average CRPS is thus mostly a measure of resolution only.




Figure 7 – CRPS averaged over space and time (left) and reliability term (right), versus the proportion of flat rank histograms, for each expert, by lead time.

References

- Adjakossa, E., Goude, Y., and Wintenberger, O. (2020). Kalman Recursions Aggregated Online. *arXiv preprint arXiv:2002.12173*.
- Allard, D., Comunian, A., and Renard, P. (2012). Probability aggregation methods in geoscience. *Mathematical Geosciences*, 44(5):545–581.
- Anderson, J. L. (1996). A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *Journal of Climate*, 9(7):1518–1530.
- Baran, S. and Lerch, S. (2015). Log-normal distribution based Ensemble Model Output Statistics models for probabilistic wind-speed forecasting. *Quarterly Journal of the Royal Meteorological Society*.
- Baran, S. and Lerch, S. (2016). Mixture EMOS model for calibrating ensemble forecasts of wind speed. *Environmetrics*, 27:116–130.
- Baudin, P. (2015). *Prévision séquentielle par agrégation d'ensemble: application à des prévisions météorologiques assorties d'incertitudes*. PhD thesis, Université Paris-Saclay.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.
- Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, 29(4):1165–1188.
- Bogner, K., Liechti, K., and Zappa, M. (2017). Combining quantile forecasts and predictive distributions of streamflows. *Hydrology and Earth System Sciences*, 21(11):5493–5502.
- Bougeault, P., Toth, Z., Bishop, C., Brown, B., Burridge, D., Chen, D. H., Ebert, B., Fuentes, M., Hamill, T. M., Mylne, K., et al. (2010). The THORPEX interactive grand global ensemble. *Bulletin of the American Meteorological Society*, 91(8):1059.
- Bröcker, J. (2009). Reliability, sufficiency, and the decomposition of proper scores. *Quarterly Journal of the Royal Meteorological Society*, 135(643):1512–1519.
- Bröcker, J. (2018). Assessing the reliability of ensemble forecasting systems under serial dependence. *Quarterly Journal of the Royal Meteorological Society*, 144(717):2666–2675.
- Buizza, R., Houtekamer, P., Pellerin, G., Toth, Z., Zhu, Y., and Wei, M. (2005). A comparison of the ECMWF, MSC, and NCEP global ensemble prediction systems. *Monthly Weather Review*, 133(5):1076–1097.

- Cesa-Bianchi, N., Lugosi, G., et al. (2006). *Prediction, learning, and games*, volume 1. Cambridge University Press Cambridge.
- Collet, J. and Richard, M. (2017). A Generic Method for Density Forecasts Recalibration. In *Forecasting and Risk Management for Renewable Energy*, pages 147–166. Springer.
- Descamps, L., Labadie, C., and Bazile, E. (2011). Representing model uncertainty using the multiparametrization method. In *Proceedings of ECMWF Workshop on Representing Model Uncertainty and Error in Numerical Weather and Climate Prediction Models, 20-24 June 2011*, pages 175–182.
- Descamps, L., Labadie, C., Joly, A., Bazile, E., Arbogast, P., and Cébron, P. (2014). PEARP, the Météo-France short-range ensemble prediction system. *Quarterly Journal of the Royal Meteorological Society*, 141(690):1671–1685.
- Elmore, K. L. (2005). Alternatives to the chi-square test for evaluating rank histograms from ensemble forecasts. *Weather and forecasting*, 20(5):789–795.
- Gerchinovitz, S., Mallet, V., and Stoltz, G. (2008). A further look at sequential aggregation rules for ozone ensemble forecasting. *Rapport technique, INRIA Paris-Rocquencourt et École normale supérieure, Paris*.
- Gneiting, T., Balabdaoui, F., and Raftery, A. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2):243–268.
- Gneiting, T., Raftery, A., Westveld III, A., and Goldman, T. (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133(5):1098–1118.
- Gneiting, T., Ranjan, R., et al. (2013). Combining predictive distributions. *Electronic Journal of Statistics*, 7:1747–1782.
- Hamill, T. (2001). Interpretation of rank histograms for verifying ensemble forecasts. *Monthly Weather Review*, 129(3):550–560.
- Hamill, T. and Colucci, S. (1998). Evaluation of Eta-RSM ensemble probabilistic precipitation forecasts. *Monthly Weather Review*, 126(3):711–724.
- Hamill, T. M. and Colucci, S. J. (1996). Random and systematic error in NMC’s short-range Eta ensembles. In *Preprints, 13th Conf. on Probability and Statistics in the Atmospheric Sciences, San Francisco, CA, Amer. Meteor. Soc.*, pages 51–56.
- Hemri, S., Scheuerer, M., Pappenberger, F., Bogner, K., and Haiden, T. (2014). Trends in the predictive performance of raw ensemble weather forecasts. *Geophysical Research Letters*, 41(24):9197–9205.
- Hersbach, H. (2000). Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, 15(5):559–570.

- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30.
- Holton, J. R. and Hakim, G. J. (2012). *An introduction to dynamic meteorology*, volume 88. Academic press.
- Jolliffe, I. and Stephenson, D. (2011). *Forecast verification: a practitioner’s guide in atmospheric science, second edition*. Wiley-Blackwell.
- Jolliffe, I. T. and Primo, C. (2008). Evaluating rank histograms using decompositions of the chi-square test statistic. *Monthly Weather Review*, 136(6):2133–2139.
- Leutbecher, M. and Palmer, T. N. (2008). Ensemble forecasting. *Journal of Computational Physics*, 227(7):3515–3539.
- Malardel, S. (2005). *Fondamentaux de météorologie: à l’école du temps*, volume 45. Cepadues.
- Mallet, V., Mauricette, B., and Stoltz, G. (2007). Description of Sequential Aggregation Methods and their Performances for Ozone Ensemble Forecasting. *Technical report, École normale supérieure, DMA and CEREAs*.
- Matheson, J. E. and Winkler, R. L. (1976). Scoring rules for continuous probability distributions. *Management science*, 22(10):1087–1096.
- Meinshausen, N. (2006). Quantile regression forests. *The Journal of Machine Learning Research*, 7:983–999.
- Möller, A. and Groß, J. (2016). Probabilistic temperature forecasting based on an ensemble autoregressive modification. *Quarterly Journal of the Royal Meteorological Society*, 142(696):1385–1394.
- Möller, D. and Scheuerer, M. (2013). *Postprocessing of Ensemble Forecasts for Wind Speed over Germany*. PhD thesis, Diploma thesis, Faculty of Mathematics and Computer Science, Heidelberg University. Available online at <http://www.rzuser.uni-heidelberg.de/~kd4/files/Moeller2013.pdf>.
- Murphy, A. (1993). What is a good forecast? an essay on the nature of goodness in weather forecasting. *Weather and Forecasting*, 8(2):281–293.
- Mylne, K. R. (2002). Decision-making from probability forecasts based on forecast value. *Meteorological Applications*, 9(3):307–315.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Richardson, D. S. (2001). Measures of skill and value of ensemble prediction systems, their interrelationship and the effect of ensemble size. *Quarterly Journal of the Royal Meteorological Society*, 127(577):2473–2489.

- Scheuerer, M., Möller, D., et al. (2015). Probabilistic wind speed forecasting on a grid based on ensemble model output statistics. *The Annals of Applied Statistics*, 9(3):1328–1349.
- Siebert, S. (2015). *SpecsVerification: Forecast Verification Routines for the SPECS FP7 Project*. R package version 0.4-1.
- Stoltz, G. (2010). Agrégation séquentielle de prédicteurs: méthodologie générale et applications à la prévision de la qualité de l’air et à celle de la consommation électrique. *Journal de la Société Française de Statistique*, 151(2):66–106.
- Swinbank, R., Kyouda, M., Buchanan, P., Froude, L., Hamill, T. M., Hewson, T. D., Keller, J. H., Matsueda, M., Methven, J., Pappenberger, F., et al. (2016). The TIGGE project and its achievements. *Bulletin of the American Meteorological Society*, 97(1):49–67.
- Taillardat, M., Mestre, O., Zamo, M., and Naveau, P. (2016). Calibrated Ensemble Forecasts using Quantile Regression Forests and Ensemble Model Output Statistics. *Monthly Weather Review*, 144(6):2375–2393.
- Talagrand, O., Vautard, R., and Strauss, B. (1997). Evaluation of probabilistic prediction systems. In *Proc. ECMWF Workshop on Predictability*, pages 1–25.
- Thorarinsdottir, T. L. and Gneiting, T. (2010). Probabilistic forecasts of wind speed: ensemble model output statistics by using heteroscedastic censored regression. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 173(2):371–388.
- Thorey, J. (2017). *Prévision d’ensemble par agrégation séquentielle appliquée à la prévision de production d’énergie photovoltaïque*. PhD thesis, INRIA. to be submitted (in English).
- Wasserstein, R., Schirm, A., and N., L. (2019). [Special Issue]Statistical Inference in the 21st Century: A World Beyond $p < 0.05$. *The American Statistician*, 73(1).
- Wilks, D. S. (2018). Enforcing calibration in ensemble postprocessing. *Quarterly Journal of the Royal Meteorological Society*, 144(710):76–84.
- Wilson, L. J., Bearegard, S., Raftery, A. E., and Verret, R. (2007). Calibrated surface temperature forecasts from the Canadian ensemble prediction system using Bayesian model averaging. *Monthly Weather Review*, 135(4):1364–1385.
- Winkler, R., Muñoz, J., Cervera, J., Bernardo, J., Blattenberger, G., Kadane, J., Lindley, D., Murphy, A., Oliver, R., and Ríos-Insua, D. (1996). Scoring rules and the evaluation of probabilities. *Test*, 5(1):1–60.
- Zamo, M., Bel, L., Mestre, O., and Stein, J. (2016). Improved gridded wind-speed forecasts by statistical post-processing of numerical models with block regression. *Weather and Forecasting*, 31(6):1929–1945.

- Zamo, M., Mestre, O., Arbogast, P., and Pannekoucke, O. (2014). A benchmark of statistical regression methods for short-term forecasting of photovoltaic electricity production. Part II: Probabilistic forecast of daily production. *Solar Energy*, 105:804–816.
- Zamo, M. and Naveau, P. (2018). Estimation of the Continuous Ranked Probability Score with Limited Information and Applications to Ensemble Weather Forecasts. *Mathematical Geosciences*, 50(2):209–234.
- Zhu, Y., Toth, Z., Wobus, R., Richardson, D., and Mylne, K. (2002). The economic value of ensemble-based weather forecasts. *Bulletin of the American Meteorological Society*, 83(1):73.