

A possibilistic interpretation of ensemble forecasts: Experiments on the imperfect Lorenz 96 system

Noémie Le Carrer¹ and Peter L. Green¹

¹Institute for Risk and Uncertainty, Department of Engineering, University of Liverpool, United Kingdom

Correspondence: N. Le Carrer (nlc@liverpool.ac.uk)

Abstract. Ensemble forecasting has gained popularity in the field of numerical medium-range weather prediction as a means of handling the limitations inherent to predicting the behaviour of high dimensional, nonlinear systems, that have high sensitivity to initial conditions. Through small strategical perturbations of the initial conditions, and in some cases, stochastic parameterization schemes of the atmosphere-ocean dynamical equations, ensemble forecasting allows one to sample possible future scenarii in a Monte-Carlo like approximation. Results are generally interpreted in a probabilistic way by building a predictive density function from the ensemble of weather forecasts. However, such a probabilistic interpretation is regularly criticized for not being reliable, because of the chaotic nature of the dynamics of the atmospheric system as well as the fact that the ensembles of forecasts are not, in reality, produced in a probabilistic manner. To address these limitations, we propose a novel approach: a possibilistic interpretation of ensemble predictions, taking inspiration from fuzzy and possibility theories. Our approach is tested on an imperfect version of the Lorenz 96 model and results are compared against those given by a standard probabilistic ensemble dressing. The possibilistic framework reproduces (ROC curve, resolution) or improves (ignorance, sharpness, reliability) the performance metrics of a standard univariate probabilistic framework.

1 Introduction

As a result of its chaotic dynamics, the prediction of the atmospheric system is particularly sensitive to the limited resolution in the initial conditions (ICs), discrepancies introduced by measurement error, computational truncation and an incomplete description of the system's dynamics (closure problem). Ensemble predictions have consequently been developed to characterize the skill of single numerical predictions of the future state of the atmosphere. As suggested by Leith (1974), assuming that the error field is dominated by observational error (i.e. error on the ICs propagated forward in the model), we can perturb M times the best estimate for the ICs, run forward the model from each IC and interpret the M results in a Monte-Carlo like fashion. In other words, we use the local density of the resulting M predictions (or members) to quantify the plausibility of a given future scenario. Instead of the traditional point deterministic predictions, probabilistic predictions are thus realized. Today, the ICs are perturbed according to various schemes, designed to sample in a minimalist way systems of millions of dimensions (like numerical weather global models). These schemes generally select the initial perturbations leading to the fastest growing perturbations (e.g. singular vectors (Hartmann et al., 1995)).

25 Yet, in practice, the assumption of a near-perfect model, where observational error is more significant than model error, is not always true. Thus, individual member trajectories are not expected to stay in the convex hull of the ensemble after a few hours (Toth and Kalnay, 1997; Orrell, 2005). While ensemble predictions is built on the idea that the range of the ensemble provides an idea of the the possible futures and that its variance is representative of the skill of the single deterministic forecast, in practice and despite the introduction of stochastic parameterization schemes to represent model error (Buizza et al., 1999),
30 the operational ensembles are overconfident: the spread is typically too small (Wilks and Hamill, 1995; Buizza, 2018). In particular, such probabilistic predictions are not reliable; on average, the probability derived for a given event does not equal the frequency of observation. Although ensemble-based probabilistic predictions present more skill than the climatology, they generally cannot be used as actionable probabilities, for the sheer fact that, by design (limited EPS size, biased sampling of ICs) and by context (flow-dependent regime error, strongly nonlinear system) they do not represent the true probabilities of
35 the system at hand (Legg and Mylne, 2004; Orrell, 2005; Bröcker and Smith, 2007, 2008; Smith, 2016). This is all the more true for extreme events, that, for dynamical reasons, cannot be associated to a high density of ensemble members; such events indeed result from nonlinear interactions at small scales, which cannot be reproduced in number in a limited-size EPS (Legg and Mylne, 2004).

A range of post-processing methods have been developed to tackle these limitations (Vannitsem et al., 2018). The classical
40 Bayesian model averaging (BMA; Raftery et al. (2005)) and non-homogeneous Gaussian regression (Gneiting et al., 2005) fit an optimized (sum of) parametric distribution(s) onto the ensemble of predictions. More recently, techniques involving recalibration by means of the probability integral transform (Graziani et al., 2019) or by using the actual probability of success of a given probabilistic threshold (Smith, 2016) were designed to address the lack of reliability. Allen et al. (2019) introduced a regime-dependent adaptation of the traditional post-processing parametric methods, to tackle the issue of possibly significant
45 model error. All ensemble post-processing techniques are trained on an archive set of (ensemble, observation) pairs, using the same model. Most often, the objective function to optimize is a performance score, like the negative log-likelihood or the continuous ranked probability score, whose individual results for each couple (ensemble, verification) are aggregated over the whole archive.

However, if post-processing globally improves the skill of common events, it deteriorates the results for extreme events
50 (Mylne et al., 2002), which generally, for predictability reasons, cannot be associated to a high density of ensemble members (Legg and Mylne, 2004).

For all these reasons, we may wonder, echoing Bröcker and Smith (2008), *whether the probability distribution (PDF) is the best representation of the valuable information contained in an EPS*. Rather, the description of possibility theory in Dubois et al. (2004): "a weaker theory than probability (...) also relevant in non-probabilistic settings where additivity no longer makes
55 sense and not only as a weak substitute for additive uncertainty measures" presents new opportunities, in a context where where conceptual and practical limitations restrict the applicability of a density-based (i.e. additive) interpretation of EPS.

This is what we investigate in this work. Namely: can we design a possibilistic framework for interpreting EPS that would perform at least as well as a standard probabilistic approach for most of the performance metrics, and improve the known shortcomings of the probabilistic approach? We investigate this question by means of numerical experiments on a commonly-

60 used surrogate model of the atmospheric dynamics, the Lorenz 96 system. Section 2 introduces the basics of possibility theory, then used to develop an original possibilistic framework for the interpretation of ensemble of predictions in Sec. 3. This framework is tested on the imperfect Lorenz 96 model in Sec. 4. A conclusion follows.

2 Possibility theory

Possibility theory is an uncertainty theory developed by Zadeh (1978) from fuzzy set theory. It is designed to handle in-
 65 complete information and represent ignorance. Considering a system whose state is described by a variable $x \in \mathcal{X}$, the possibility distribution $\pi : \mathcal{X} \mapsto [0, 1]$ represents the state of knowledge of an agent about the current state of the system. Given an event $A = \{x \in S_A\}$, the possibility and necessity measures are defined respectively as: $\Pi(A) = \sup_{x \in S_A} \pi(x)$ and $N(A) = 1 - \Pi(\bar{A})$ where \bar{A} represents the complementary event of A . $\Pi(A)$ and $N(A)$ satisfy the following axioms:

1. $\Pi(\mathcal{X}) = 1$ and $\Pi(\emptyset) = 0$
- 70 2. $\Pi(A \cup B) = \max(\Pi(A), \Pi(B))$ (similar to $N(A \cap B) = \min(N(A), N(B))$), where $B = \{x \in S_B\}$.

The following conventions apply (Cayrac et al., 1994):

- a. $N(A) = 1 \Leftrightarrow \Pi(\bar{A}) = 0$ indicates that A has to happen, it is necessary;
- b. $0 < N(A) < 1$ is a tentative acceptance of A to a degree $N(A)$, since $\min(N(A), N(\bar{A})) = 0$ from axiom 2 (\bar{A} is not necessary at all);
- 75 c. $(\Pi(A) = \Pi(\bar{A}) = 1) \Leftrightarrow (N(A) = N(\bar{A}) = 0)$ represents total ignorance: the evidence doesn't allow us to conclude if A is rather true or false;

Possibility and probability distributions are interconnected, through the description of uncertainty by imprecise probabilities (c.f. the Dempster-Shafer Theory of Evidence's framework (Dempster, 2008)). Under specific constraints, an imprecise distribution can degenerate into either a probability or a possibility distribution. One can consequently assess the degree of
 80 consistency of a possibility and a probabilistic distributions. Among the definitions of consistency (Delgado and Moral, 1987), we retain here the view of Dubois et al. (2004), that a probability measure P and possibility measure Π are consistent if the probability of all possible events A satisfies $P(A) \leq \Pi(A)$. It implies, from the definition of necessity, that the probability $P(A)$ is bounded as well from below by the necessity measure: $N(A) \leq P(A) \leq \Pi(A)$. Necessity and possibility measures can consequently be viewed as upper and lower limits on the probability of a given event.

85 3 Possibilistic framework for EPS interpretation

The statistical post-processing of EPS generates forecasts in the form of predictive probability distributions $p(x|\tilde{x}, \theta)$, noted $p(x|\tilde{x})_\theta$, where $\tilde{x} = \{\tilde{x}_1, \dots, \tilde{x}_M\}$ is the ensemble, θ a vector of parameters and p a (sum of) parametric distribution(s). BMA

distributions are weighted sums of M parametric probability distributions, each one centered around a linearly corrected ensemble member. In this work, the members are exchangeable, so the mixture coefficients and parametric distributions do not vary between members and the BMA comes down to an ensemble dressing procedure. We compare our method against a Gaussian ensemble dressing, whose predictive probability distribution reads:

$$p(x|\tilde{\mathbf{x}})_\theta = \frac{1}{M} \sum_{i=1}^M \mathcal{N}(a\tilde{x}_i + \omega, \sigma^2) \quad (1)$$

where $\mathcal{N}(\mu, v)$ is the normal distribution of mean μ and variance v . The parameters $\theta = \{a, \omega, \sigma\}$ are inferred through the optimization of a performance metric, e.g. the ignorance score (Roulston and Smith, 2002), or negative log-likelihood, a strictly proper¹ and local² logarithmic score.

Here, instead of performing a probabilistic ensemble dressing, we can perform a *possibilistic* ensemble dressing: a possibilistic membership function is dressed around each ensemble member first shifted and scaled. Similarly to its probabilistic twin, the i th possibility kernel is assumed to represent the possibility distribution of the true state of the system, given the observation of \tilde{x}_i . Because we have several member observations $i = \{1, \dots, M\}$ and there is only one truth (the actual system's state), we can interpret it as a union (OR) of possibilities. Fuzzy set theory offers several definitions for computing the distribution resulting of the union of two fuzzy distributions. We adopt here the max-sum definition: $\pi_{A \cup B}(x) = \max(\pi_A(x), \pi_B(x))$, although some of our tests, not presented here, show that alternative definitions do not significantly change results.

Gaussian kernels $\exp^{-\frac{1}{2}u_i^2}$ are thus fitted to each member \tilde{x}_i , with $u_i = \frac{x - (a\tilde{x}_i + \omega)}{\sigma}$, a the scaling factor, ω the shifting of the kernels' peaks from the individual member \tilde{x}_i and σ a parameter accounting for the width of the individual kernels. The resulting possibilistic distribution is given by the sum, in a possibilistic manner, of all the individual kernels:

$$\pi(x) = \cup_{i=1 \dots M} \exp^{-\frac{(x - (a\tilde{x}_i + \omega))^2}{2\sigma^2}} = \sup_{i=1 \dots M} \exp^{-\frac{(x - (a\tilde{x}_i + \omega))^2}{2\sigma^2}} \quad (2)$$

For any event of interest $A = \{x \in S_A\}$, we can extract the possibility and necessity measures $\Pi(A, \theta)$ and $N(A, \theta)$ (noted $\Pi_\theta(A)$ and $N_\theta(A)$), given the knowledge encoded in $\pi(x, \theta)$ (noted π_θ). $\Pi_\theta(A)$ evaluates to what extent A is logically consistent with π_θ whereas $N_\theta(A)$ evaluates to what extent A is certainly implied by π_θ . Ideally, this pair falls in an area of the possibilistic diagram (N, Π) that is close to one of the three notable points: $(1, 1)$ for A certain; $(0, 0)$ for \bar{A} certain; $(0, 1)$ for total ignorance, i.e. both A and \bar{A} are possible but none is necessary given π . Points on the line $N = 0$ are in favor of \bar{A} , the more favorable the closer to $(0, 0)$; points on the line $\Pi = 1$ are in favor of A , the more favorable the closer to $(1, 1)$. Other areas of the diagram are inconsistent with the axioms defining Π and N .

From the geometric interpretation given by the possibilistic diagram, several options are available for scoring each point $(N_\theta(A), \Pi_\theta(A))$ that is, for assessing the quality of the prediction given by the pair $(N_\theta(A), \Pi_\theta(A))$. A brute-force method is to minimize the distance to the correct pole (e.g. $(1, 1)$ for A true). Yet, such an approach would try and push events towards $(1, 1)$ or $(0, 0)$ on the possibilistic diagram, thus ignoring the ignorance pole and, as a result, the idea that some events are impossible to predict from a particular EPS set. A more complete method could, for instance, also consider the rank r of the

¹i.e. it takes its optimal value only when the forecast probability is equal to the true distribution of the system.

²i.e. it does not depend on the full forecast distribution, but only on the predictive probability associated to the true system's state.

EPS w.r.t. A . Namely, if the actual observation x^* is in S_A , the associated point should belong to the line $\Pi = 1$ but the distance
 120 to the ignorance pole $(1, 0)$ should be proportional to r . The same applies for $x^* \notin S_A$; the associated point should belong to
 line $N = 0$ with the distance to $(1, 0)$ proportional to $r_{\bar{A}} = M - r_A$. Thus, an observation $x^* \in S_A$ associated to an erroneous
 EPS ($r \rightarrow 0$) will fall close to the ignorance pole, suggesting that we cannot trust the raw ensemble. A score verifying these
 requirements is:

$$S_i(\theta) = \begin{cases} |N_\theta(A) - \frac{r}{M}| + |\Pi_\theta(A) - 1|, x^* \in S_A \\ N_\theta(A) + |\Pi_\theta(A) - \frac{r}{M}|, x^* \notin S_A \end{cases}$$

125 Given a training set containing n pairs (\tilde{x}_i, x_i^*) , the final empirical score is: $S(\theta) = \frac{1}{n} \sum_{i=1}^n S_i(\theta)$ and training consists of
 finding the θ that minimizes S .

4 Application to the imperfect Lorenz 96 system

To test our framework, we reproduce the experiment designed by Williams et al. (2014), who used an imperfect L96 model
 (Lorenz, 1996) to generate ensemble predictions and investigate the performance of ensemble post-processing methods for the
 130 prediction of extreme events. The training sets consist of 4000 independent pairs of EPS of size $M = 12$ and the associated
 observations, for each lead time $\tau = \{1, 3, 5, 7\}$ days³. The EPS have beforehand been pre-processed to remove the constant
 bias. The testing set consists of another 10,000 independent pairs of bias-corrected EPS and associated observations, for each
 lead time. We consider the prediction of an extreme event: $A_e = \{x \leq q_{0.05}\}$, where $q_{0.05}$ is the 0.05 quantile of the climatic
 distribution of x and a common event $A_c = \{q_{0.5} \leq x \leq q_{0.6}\}$. Results are compared against those given by a probabilistic
 135 post-processing, namely a Gaussian ensemble dressing.

We first assess the performance of each interpretation in terms of the empirical ignorance score relative to the climatology:

$$S_n(p_\theta, c) = \frac{1}{n} \sum_{i=1}^n \left(IGN(r_\theta, x_i^*) - IGN(c, x_i^*) \right) = -\frac{1}{n} \sum_{i=1}^n \log_2 \left(\frac{r_\theta(x_i^*)}{c(x_i^*)} \right) \quad (3)$$

where, following the work of Bröcker and Smith (2008), in the probabilistic framework, the predictive probability $p_\theta(x^*|\tilde{x})$ is
 blended with the climatology $c(x^*)$ of the verification x^* : $r_\theta(x^*) = \alpha p_\theta(x^*) + (1 - \alpha)c(x^*)$. Our possibilistic framework is
 140 a mapping $\mathbb{R}^M \mapsto [0, 1] \times [0, 1]$, while the ignorance applies to a probabilistic prediction $\mathbb{R}^M \mapsto [0, 1]$. We consequently need
 to find a mapping from the dual measures N and Π to an equivalent probability. Since possibility and necessity measures can
 be seen as upper and lower bounds of a consistent probability measure, we can write $P(A) = \alpha N(A) + (1 - \alpha)\Pi(A)$ with
 $\alpha \in [0, 1]$ for any event A of interest. Varying α allows one to browse across the range of associated probabilities, consistent
 with the possibility distribution π . We use this technique to compute the ignorance score of the possibilistic framework and
 145 compare its range to the performance of a probabilistic Gaussian ensemble dressing. Both frameworks are characterized by
 negative relative ignorance, confirming that they have a predictive added-value over climatology. The difference in ignorance
 equals the difference in expected returns that one would get by placing bets proportional to their probabilistic forecasts.

³ $\tau = 1$ corresponds to 0.2 model time units after initialization and can be associated with approximately 1 day in the real world (Lorenz, 1996).

As shown in Figure 1, for both types of events, the possibilistic framework performs as well or slightly better than the probabilistic, for all $\alpha \in [0, 1]$. The slight increase in performance remains relatively constant or even improve (extreme event case) with lead time. The relative ignorance of the possibilistic framework has a variance (due to the range of α) that grows with the lead time, as expected.

To understand better the operational consequences of such results, we report in Figure 2 the relative operating characteristic (ROC) of both frameworks at lead times of 3 and 7 days. Given a binary prediction (yes/no w.r.t. event A), the ROC plots the hit rate (HR; fraction of correctly predicted A over all A observed) versus the false alarm rate (FA; fraction of wrongly predicted A over all \bar{A} observed). We use increasing thresholds $p_t \in [0, 1]$ for making the decision (yes if $P(A) \geq p_t$) and report the associated HR and FA in the graph. Again, we vary α to see the range of HR and FA covered for each p_t by the possibilistic prediction (N, Π). The resulting points form a curve (probabilistic approach) or a cloud (possibilistic method), which are a visual way to assess the ability of a forecast system to discriminate between events and non events.

The possibilistic curves all fit or are very close to the probabilistic curves, for both extreme and common events and for all lead times. The main difference is their extension: the possibilistic framework remains located in areas of relatively small FA, compared to the results of the probabilistic approach for similar thresholds p_t . This results indicates that the HR remains smaller than what can be achieved by the probabilistic framework, showing lower skill. The fact that the possibilistic curves yet lies on the probabilistic ROC curves shows that the reason behind this discrepancy is not a lack of discrimination between events and non-events; for a given FA, both methods provide the same HR. The reason is connected to a bias in probabilities for the possibilistic approach towards zero and towards 1: the possibilistic framework is very sharp, as shown on the diagrams in Figure 3. Because they are not blended with climatology, a large part of the predictions have zero probability associated to the event of interest, instead of a minimal one, which prevents the current implementation of the possibilistic framework from reaching higher HR. Side experimentation not reproduced here has shown that weighting the scores attributed to observed event A in the global empirical training score allows to reproduce fully the probabilistic curve for each lead time.

Reliability diagrams presented on Figure 4 plot the observed conditional frequencies against the corresponding forecast probabilities for lead time 3 and 7 days. They illustrate how well the predicted probabilities of an event correspond to their observed conditional frequencies. The predictive model is all the more reliable (i.e. actionable) when the associated curve is close to the diagonal. Noting that the diagonal represents perfect reliability, the distance to the diagonal indicates underforecasting (curves above) or overforecasting (curves below). Distance above the horizontal climatology line indicates a system with resolution, a system that does discriminate between events and non-events. The cones defined by the no-skill line (half-way between the climatology and perfect reliability) and the vertical climatology line allow us to define areas where the forecast system is skilled.

The probabilistic curves are globally aligned with the perfect reliability line, yet with growing lead time, they are restricted to small probabilities only (because of wider EPS or pure predictability issues such as mentioned for extreme events). On the contrary, the reliability plots associated with the possibilistic approach cover all range of probabilities. This approach tends to be underforecasting (resp. overforecasting) for small (resp. large) probabilities, especially for the common event. A large part of the area covered by the possibilistic solutions is contained in the skill cones for the rare event, denoting a skilled predictive

system for all but very low predictive probabilities. Results are less interesting for the common event, where the possibilistic framework leads to a flatter diagram, indicating less resolution, especially with larger lead times.

185 **5 Conclusions**

In this work, we have presented a possibilistic framework which allows us to interpret ensemble predictions without the notion of *member density*, or *additivity* that proved to be incoherent with the conditions in which EPS were built. Preliminary results show that such a framework can be used to reproduce the probabilistic performances (ROC curves, resolution) and even slightly improve some of them (ignorance, sharpness, reliability). Moreover, the proposed approach addresses some of the well-known
190 limitations of the probabilistic framework (reliability, for example). The added-value of this framework is particularly tangible for extreme events. Further work is needed to improve the design of the possibilistic distributions, by means of dynamical information or statistical priors. Besides, developments regarding the understanding and the operational use of such 'fuzzy' results are necessary.

Author contributions. NLC conceived of the presented idea, designed and implemented the research and wrote the article. PG contributed to
195 the analysis of the results and reviewed the article.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This research is funded by the Engineering & Physical Sciences Research Council (EPSRC) and the Economic & Social Research Council (ESRC) with grant no. EP/L015927/1.

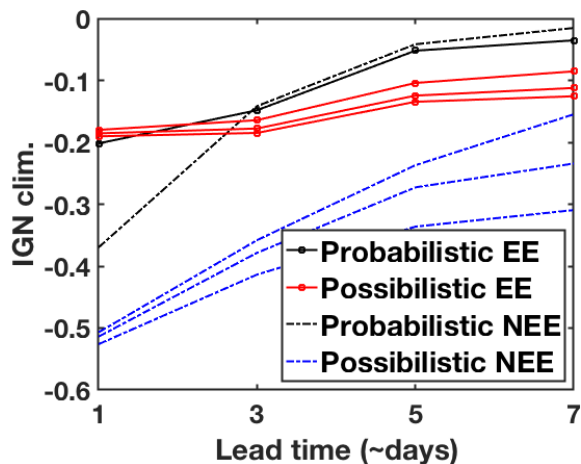


Figure 1. Ignorance relative to the climatology computed for the possibilistic (colored lines) and probabilistic (black lines) frameworks, in the case of the prediction of an extreme (EE; solid line) and a common (NEE; dashed line) event of interest, as defined in Sec. 4. The upper and lower bounds, as well as the median, obtained by considering that $N(A) \leq P(A) \leq \Pi(A)$ in the possibilistic framework are reported.

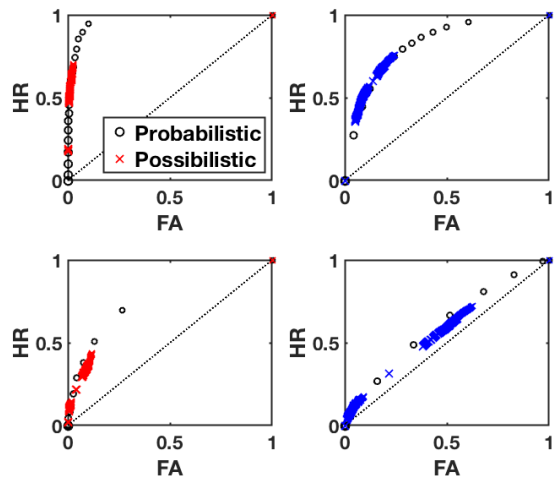


Figure 2. ROC curves for the extreme event (left side) and common event (right side) at lead time 3 days (top) and 7 days (bottom). The probabilistic results are reported by means of black circles and the possibilistic results by means of colored crosses. The larger the symbol, the larger the threshold probability used to compute HR and FA.

References

200 Allen, S., Ferro, C. A., and Kwasniok, F.: Regime-dependent statistical post-processing of ensemble forecasts, Quarterly Journal of the Royal Meteorological Society, 145, 3535–3552, 2019.




Figure 3. Normalized histograms of the equivalent forecast probabilities in the possibilistic framework for the observations of the extreme event (left) and common event (right) at lead time 3 days. The corresponding distributions of predictive probabilities in the probabilistic framework come on top as a thick black lines.




Figure 4. Reliability diagrams for the extreme event (left side) and common event (right side) at lead time 3 (top) and 7 days (bottom). The probabilistic results are reported in black line, while the upper, median and lower bounds of the possibilistic ones are in thinner red lines. Standards elements of comparison are reported in the diagram, as described in Sec. 4, namely the diagonal (perfect reliability), the climatological reference (horizontal dotted) and the cones of skill (inside the dashed-dotted secants).

Bröcker, J. and Smith, L. A.: Increasing the reliability of reliability diagrams, *Weather and forecasting*, 22, 651–661, 2007.

Bröcker, J. and Smith, L. A.: From ensemble forecasts to predictive distribution functions, *Tellus A: Dynamic Meteorology and Oceanography*, 60, 663–678, 2008.

205 Buizza, R.: Ensemble forecasting and the need for calibration, in: *Statistical Postprocessing of Ensemble Forecasts*, pp. 15–48, Elsevier, 2018.

Buizza, R., Milleer, M., and Palmer, T. N.: Stochastic representation of model uncertainties in the ECMWF ensemble prediction system, *Quarterly Journal of the Royal Meteorological Society*, 125, 2887–2908, 1999.

- Cayrac, D., Dubois, D., Haziza, M., and Prade, H.: Possibility theory in " Fault mode effect analyses". A satellite fault diagnosis application, 210 in: Proceedings of 1994 IEEE 3rd International Fuzzy Systems Conference, pp. 1176–1181, IEEE, 1994.
- Delgado, M. and Moral, S.: On the concept of possibility-probability consistency, *Fuzzy sets and Systems*, 21, 311–318, 1987.
- Dempster, A. P.: Upper and lower probabilities induced by a multivalued mapping, in: *Classic Works of the Dempster-Shafer Theory of Belief Functions*, pp. 57–72, Springer, 2008.
- Dubois, D., Foulloy, L., Mauris, G., and Prade, H.: Probability-possibility transformations, triangular fuzzy sets, and probabilistic inequalities, 215 *Reliable computing*, 10, 273–297, 2004.
- Gneiting, T., Raftery, A. E., Westveld III, A. H., and Goldman, T.: Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation, *Monthly Weather Review*, 133, 1098–1118, 2005.
- Graziani, C., Rosner, R., Adams, J. M., and Machete, R. L.: Probabilistic Recalibration of Forecasts, arXiv preprint arXiv:1904.02855, 2019.
- Hartmann, D., Buizza, R., and Palmer, T. N.: Singular vectors: The effect of spatial scale on linear growth of disturbances, *Journal of the* 220 *atmospheric sciences*, 52, 3885–3894, 1995.
- Legg, T. and Mylne, K.: Early warnings of severe weather from ensemble forecast information, *Weather and Forecasting*, 19, 891–906, 2004.
- Leith, C.: Theoretical skill of Monte Carlo forecasts, *Monthly Weather Review*, 102, 409–418, 1974.
- Lorenz, E. N.: Predictability: A problem partly solved, in: *Proc. Seminar on predictability*, vol. 1, 1996.
- Mylne, K., Woolcock, C., Denholm-Price, J., and Darvell, R.: Operational calibrated probability forecasts from the ECMWF ensemble 225 prediction system: implementation and verification, in: *Preprints of the Symposium on Observations, Data Assimilation and Probabilistic Prediction*, pp. 113–118, 2002.
- Orrell, D.: Ensemble forecasting in a system with model error, *Journal of the atmospheric sciences*, 62, 1652–1659, 2005.
- Raftery, A. E., Gneiting, T., Balabdaoui, F., and Polakowski, M.: Using Bayesian model averaging to calibrate forecast ensembles, *Monthly weather review*, 133, 1155–1174, 2005.
- 230 Roulston, M. S. and Smith, L. A.: Evaluating probabilistic forecasts using information theory, *Monthly Weather Review*, 130, 1653–1660, 2002.
- Smith, L. A.: Integrating information, misinformation and desire: improved weather-risk management for the energy sector, in: *UK Success Stories in Industrial Mathematics*, pp. 289–296, Springer, 2016.
- Toth, Z. and Kalnay, E.: Ensemble forecasting at NCEP and the breeding method, *Monthly Weather Review*, 125, 3297–3319, 1997.
- 235 Vannitsem, S., Wilks, D. S., and Messner, J.: *Statistical postprocessing of ensemble forecasts*, Elsevier, 2018.
- Wilks, D. S. and Hamill, T. M.: Potential economic value of ensemble-based surface weather forecasts, *Monthly weather review*, 123, 3565–3575, 1995.
- Williams, R., Ferro, C., and Kwasniok, F.: A comparison of ensemble post-processing methods for extreme events, *Quarterly Journal of the Royal Meteorological Society*, 140, 1112–1120, 2014.
- 240 Zadeh, L. A.: Fuzzy sets as a basis for a theory of possibility, *Fuzzy sets and systems*, 1, 3–28, 1978.