

Locally tail-scale invariant scoring rules for evaluation of extreme value forecasts

Downloaded from: https://research.chalmers.se, 2024-07-27 03:12 UTC

Citation for the original published paper (version of record):

Ólafsdóttir, H., Rootzen, H., Bolin, D. (2024). Locally tail-scale invariant scoring rules for evaluation of extreme value forecasts. International Journal of Forecasting, In Press. http://dx.doi.org/10.1016/j.ijforecast.2024.02.007

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

International Journal of Forecasting xxx (xxxx) xxx

Contents lists available at ScienceDirect



International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast



Locally tail-scale invariant scoring rules for evaluation of extreme value forecasts

Helga Kristin Olafsdottir^{a,*}, Holger Rootzén^a, David Bolin^b

^a Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, Gothenburg, Sweden ^b Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

ARTICLE INFO

Keywords: Proper scoring rules Extreme value theory CRPS SwCRPS Local tail-scale invariance

ABSTRACT

Statistical analysis of extremes can be used to predict the probability of future extreme events, such as large rainfalls or devastating windstorms. The quality of these forecasts can be measured through scoring rules. Locally scale invariant scoring rules give equal importance to the forecasts at different locations regardless of differences in the prediction uncertainty. This is a useful feature when computing average scores but can be an unnecessarily strict requirement when one is mostly concerned with extremes. We propose the concept of local weight-scale invariance, describing scoring rules fulfilling local scale invariance in a certain region of interest, and as a special case, local tailscale invariance for large events. Moreover, a new version of the weighted continuous ranked probability score (wCRPS) called the scaled wCRPS (swCRPS) that possesses this property is developed and studied. The score is a suitable alternative for scoring extreme value models over areas with a varying scale of extreme events, and we derive explicit formulas of the score for the generalised extreme value distribution. The scoring rules are compared through simulations, and their usage is illustrated by modelling extreme water levels and annual maximum rainfall, and in an application to non-extreme forecasts for the prediction of air pollution.

© 2024 The Authors. Published by Elsevier B.V. on behalf of International Institute of Forecasters. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

The main aim of the statistical analysis of extremes in fields such as hydrology, meteorology, climate science, and finance is to forecast, or predict, the risk of the occurrence of future dangerous extreme events, such as huge rainfall events or very large fluctuations in the prices of a financial portfolio. Typically, the forecast is in terms of an estimated distribution, for example the estimated distribution of the size of the largest event which will occur during some specified period of time. In this case,

* Corresponding author.

the forecast is said to be probabilistic. The forecast distribution is then often presented to decision makers in a reduced form, perhaps as a predicted probability of exceeding one or several high thresholds. In many cases, there are several competing forecast distributions, and a standard way of choosing between these competing forecasts is to compute goodness-of-fit measures of the distributions based on the existing data, and then use the distribution that fits best. This approach is, for example, broadly used in hydrology (Cugerone & De Michele, 2015). However, an alternative gaining increasing interest is to base the selection on proper scoring rules (Zamo & Naveau, 2018).

A scoring rule is a functional $S(\mathbb{P}, y)$ that takes in an outcome y and a probabilistic forecast \mathbb{P} , the predictive distribution for the outcome. If the outcome $Y \sim \mathbb{Q}$ is a

https://doi.org/10.1016/j.ijforecast.2024.02.007

0169-2070/© 2024 The Authors. Published by Elsevier B.V. on behalf of International Institute of Forecasters. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Please cite this article as: H.K. Olafsdottir, H. Rootzén and D. Bolin, Locally tail-scale invariant scoring rules for evaluation of extreme value forecasts. International Journal of Forecasting (2024), https://doi.org/10.1016/j.ijforecast.2024.02.007.

E-mail addresses: khelga@chalmers.se (H.K. Olafsdottir), hrootzen@chalmers.se (H. Rootzén), david.bolin@kaust.edu.sa (D. Bolin).

H.K. Olafsdottir, H. Rootzén and D. Bolin

random variable with distribution \mathbb{Q} , the score $S(\mathbb{P}, Y)$ is also a random variable, with the expected value denoted by $S(\mathbb{P}, \mathbb{Q})$. To be able to make an earnest prediction, it is preferred that the expected score is maximised under the correct model, i.e. that $S(\mathbb{Q}, \mathbb{Q}) \geq S(\mathbb{P}, \mathbb{Q})$, for all predictive distributions \mathbb{P} . In this case the score is said to be proper (Gneiting & Raftery, 2007). If strict inequality holds for all $\mathbb{P} \neq \mathbb{Q}$, the scoring rule is strictly proper.

An example of a commonly used proper scoring rule is the continuous ranked probability score (CRPS), which is defined as

$$\operatorname{CRPS}(\mathbb{P}, y) = -\int_{\mathbb{R}} (F(x) - 1(x \ge y))^2 dx,$$

where *F* is the cumulative distribution function of the forecast distribution \mathbb{P} , and $1(\cdot)$ is the indicator function. Another popular example is the logarithmic score $LS(\mathbb{P}, y) = \log f(y)$, where f is the density function of P. Both of these scores are widely used in meteorology, climate science, and finance (Haiden et al., 2019; Ingebrigtsen et al., 2015; Opschoor et al., 2017). The choice of scoring rule allows an evaluator to emphasise the desired features of the prediction. For extreme events, one typically is interested in the behaviour of the prediction in the tail of the distribution. One way of doing this is simply to multiply the proper score S_0 with a weight function wand form a new score $S(\mathbb{P}, y) = w(y)S_0(\mathbb{P}, y)$. However, this score is not proper unless w(y) is constant (Gneiting & Ranjan, 2011), and alternative ways of weighting should be used. Several such proper weighted scores exist (see e.g. Diks et al., 2011; Gneiting & Ranjan, 2011; Tödter & Ahrens, 2012). A popular one is the threshold weighted CRPS (wCRPS) by Gneiting and Ranjan (2011), where a weight function is included in the kernel of the CRPS:

wCRPS(
$$\mathbb{P}, y$$
) = $-\int_{\mathbb{R}} w(x)(F(x) - 1(x \ge y))^2 dx$.

Here w(x) is a non-negative function that can be chosen to put more emphasis on large values. However, what weight function to use is not obvious. The forecaster's dilemma, as presented by Lerch et al. (2017), describes the problem that forecasters can be encouraged to exaggerate their predictions, since the effect of missing a calamity is worse than wrongfully predicting one. To avoid this, the choice of the weight function should ideally be made by the user of the forecasts rather than by the forecaster.

In the case of multiple observations, $\mathbf{y} = (y_1, \dots, y_n)$, and a forecast $\mathbb{P} = (\mathbb{P}_1, \dots, \mathbb{P}_n)$, the joint score of the forecast is defined as the average score,

$$S(\mathbb{P}, \boldsymbol{y}) = \frac{1}{n} \sum_{i=1}^{n} S(\mathbb{P}_i, y_i),$$

and this is used to inform the selection of prediction method.

As recently shown by Bolin and Wallin (2023), many common scoring rules, such as the CRPS, are scale dependent in the sense that they put higher importance on forecasts with higher prediction uncertainty in the case when the observations y_i in the average score have varying predictability. This is sometimes an undesirable property.

For example, if one wants to predict extreme rainfalls in some spatial area where the scale of the events varies, infrastructure in the parts where the scale of events is small and in parts where the scale is large will have adapted to these differences, and accuracy is equally important throughout the area. Bolin and Wallin (2023) introduced the concept of local scale invariance to describe scoring rules which do not have this feature, and showed that the logarithmic score is an example of a score that is both strictly proper and locally scale invariant. However, the use of the logarithmic score requires computing the probability density function of the predictive distribution, which is not always easy to compute, or might not even exist.

Even though local scale invariance is an important property, it might be an unnecessarily strict requirement if one is mostly concerned with extremes. In this work, we therefore introduce a new concept-local weight-scale invariance-to describe scoring rules which have the property of local scale invariance for a certain region of interest. As a special case, when the region of interest is large events, we obtain the concept of local tail-scale invariance. We develop and study a new version of the wCRPS called the scaled wCRPS, or swCRPS, which is locally weight-scale invariant. The generalised extreme value (GEV) distribution is of special interest in extreme value statistics, and we derive explicit formulas of swCRPS for the GEV distribution. These scores are analysed in terms of local scale invariance. We compare the swCRPS and the censored likelihood score to the more common logarithmic score, the CRPS, and the wCRPS in a simulation study and in case studies where different models for water levels in the Great Lakes. extreme rainfall events in the northeastern U.S., and air pollution in the Piedmont region in Italy are evaluated using the different scoring rules.

The article is structured as follows: Section 2 gives background on scoring rules, weighting of scores, and local scale independence of scores. In Section 3, the scale independence of weighted scores is discussed. Section 4 contains simulation studies, and Section 5 presents case studies. Finally, a discussion and conclusions can be found in Section 6. The derivation of explicit formulas for the swCRPS for the GEV distribution, background on modelling and scoring of extremes, and proofs of propositions are included in the appendices of the article.

2. Background

In this section we provide more background information about the CRPS and its weighted variants, provide details about scale dependence of scoring rules, and introduce some commonly used models for annual flow and precipitation maxima. In the following, $E_{\mathbb{P}}[g(X)]$ denotes the expectation of g(X) when $X \sim \mathbb{P}$ is a random variable with distribution \mathbb{P} , and $E_{\mathbb{P},\mathbb{Q}}[g(X, Y)]$ denotes the expectation of g(X, Y) for independent random variables $X \sim \mathbb{P}$ and $Y \sim \mathbb{Q}$.

2.1. The CRPS and weighted scoring rules for extremes

A different representation of CRPS, obtained by Baringhaus and Franz (2004), is

$$\operatorname{CRPS}(\mathbb{P}, y) = \frac{1}{2} E_{\mathbb{P}, \mathbb{P}}[|X - X'|] - E_{\mathbb{P}}[|X - y|], \tag{1}$$

where X, X' are independent random variables with the same distribution \mathbb{P} and finite first moment. This expression can be derived by noting that

$$|X - Y| = \int_{\mathbb{R}} 1\{X \le u \le Y\} + 1\{Y \le u \le X\}du$$

in Eq. (1), using Fubini's theorem to compute E|X - Y|in terms of the distribution functions *F* and *G* of *X* and *Y*, respectively, and finally setting $G(x) = 1\{x \ge y\}$. This representation is useful in cases where a closed-form expression for *F* does not exist, which might be the case when using ensemble forecasts in weather prediction, or when evaluating the predictive performance of complicated hierarchical models. In these cases, the expected values can be approximated through Monte Carlo simulation from the predictive distribution. Zamo and Naveau (2018) compared different estimators of CRPS and made a recommendation on how to choose the most accurate one based on the ensemble type available.

Similarly for the wCRPS, one can show that

wCRPS(
$$\mathbb{P}, y$$
) = $\frac{1}{2} E_{\mathbb{P},\mathbb{P}}[g_w(X, X')] - E_{\mathbb{P}}[g_w(X, y)],$ (2)

where $g_w(x, x') := \left| \int_{x'}^{x} w(t) dt \right|$, provided the expectations above are finite. If one is interested only in the uppertail behaviour, one might choose the indicator weight function

$$w_u(x) \coloneqq 1(x \ge u),\tag{3}$$

for some real *u*. Instead of choosing the weight function as the indicator function w_u , one might want to consider some other variation to incorporate the whole dataset instead of just the extremes. For example, Thorarinsdottir and Schuhen (2018) suggested $w_2(y) = 1 + 1w_u(y)$ and $w_3(y) = 1 + w_u(y)u$ as alternatives, where *u* was chosen as the 97.5% observed quantile. In an example in Subsection 6.3 of Thorarinsdottir and Schuhen (2018), the best choice was to use w_2 as a weight function in the wCRPS. All these weight functions can be written as $w(y) = a + bw_u(y)$, where *a*, *b*, *u* are constants and result in a weighted sum of the wCRPS with threshold weight function and the CRPS.

Brehmer and Strokorb (2019) and Taillardat et al. (2022) show that expected scores are not suitable for scoring max-functionals, that is, functionals determined by the behaviour of the distribution at $+\infty$. An example is the so-called extreme value index. They then argue that this means that using expected scores may be unsuitable for extremes. However, standard practical use of extreme value statistics is not only concerned with the prediction of tail functionals. If one is interested in predicting maxima, e.g. yearly maxima of daily rainfall, one obtains a sample of earlier yearly maxima, fits some distribution, perhaps a GEV distribution, to this sample,

and uses the fitted distribution for prediction. This is completely parallel to fitting a normal distribution to a sample and using the fitted distribution for prediction. The understanding and use of expected scoring rules are the same for the two cases. If one instead is interested in predicting occurrences and sizes of excesses of a high threshold, one obtains a sample of such excesses and uses the sample for prediction. Again this is completely parallel to using a fitted normal distribution for prediction. Thus, we believe that there is still value in using expected scoring rules for extremes, as long as one is not solely interested in max-functionals such as the extreme value index.

Taillardat et al. (2022) also suggested treating observed scores as random variables and using qq- and pp-plots and Cramer-von Mises statistics to compare scores. This seems like a useful idea, both for scoring extremes and for scoring non-extreme values, although it is not entirely clear how to use this to obtain generally applicable rules for ranking models.

2.2. Local scale invariance and kernel scores

As mentioned above, the CPRS is a scale-dependent scoring rule. Bolin and Wallin (2023) argued that this can sometimes lead to wrong conclusions, e.g., if one uses predictions from different spatial locations, such as in the extreme rainfall analysis discussed in the introduction. In this section, we make this notion of scale dependence more specific and discuss kernel scores.

For $\theta = (\mu, \sigma) \in \mathbb{R} \times \mathbb{R}^+$ and a given probability measure \mathbb{Q} , let \mathbb{Q}_{θ} denote a location-scale transform of this measure. That is, if $Z \sim \mathbb{Q}$, then $\mu + \sigma Z \sim \mathbb{Q}_{\theta}$. Note that $\mathbb{Q}_{(0,1)} = \mathbb{Q}$. Now consider a small location-scale misspecification of θ that is proportional to σ , i.e., $\mathbb{Q}_{\theta+k\sigma r}$, where $\mathbf{r} = (r_1, r_2)$ is a two-dimensional unit vector representing the direction of perturbation, and $k \in \mathbb{R}$ is a constant. For example, a model with the correct location, μ , but a perturbation of the scale parameter $\sigma + k\sigma$, can be written as $\mathbb{Q}_{\theta+k\sigma r}$ with $\mathbf{r} = (0, 1)$. The idea of local scale invariance is that, for small perturbations, the difference between the score for the perturbed model and the true model should not depend on the scale. We make this precise in the following definition.

In the definition, and below, for a scoring rule *S* we use the notation \mathcal{P}_S for the set of probability measures on (Ω, \mathcal{F}) , such that if $\mathbb{P} \in \mathcal{P}_S$ and $\mathbb{Q} \in \mathcal{P}_S$, then $|S(\mathbb{P}, \mathbb{Q})| < \infty$. Further, for a set of probability measures \mathcal{Q}_0 , we write $\mathcal{Q} = {\mathbb{Q}}_{\theta} : \mathbb{Q} \in \mathcal{Q}_0, \theta \in \mathbb{R} \times \mathbb{R}^+$ for the set of probability measures in \mathcal{P}_S , which can be obtained as location-scale transforms of measures in \mathcal{Q}_0 . If the scoring rule *S* is proper and twice differentiable with respect to θ , a second-order Taylor expansion yields that

$$S(\mathbb{Q}_{\theta},\mathbb{Q}_{\theta}) - S(\mathbb{Q}_{\theta+k\sigma r},\mathbb{Q}_{\theta}) = k^2 \sigma^2 r^T s(\mathbb{Q}_{\theta})r + o(k^2), \quad (4)$$

as $k \searrow 0$. Since *S* is proper, the score $S(\mathbb{Q}_{\theta+k\sigma r}, \mathbb{Q}_{\theta})$ has a maximum at k = 0. Hence, $\nabla_{\theta}S(\mathbb{Q}_{\theta}, \mathbb{Q})|_{\mathbb{Q}=\mathbb{Q}_{\theta}} = 0$, and the first-order term of the Taylor expansion vanishes. Here, the function $s(\mathbb{Q}_{\theta}) := \frac{1}{2} \nabla_{\theta}^2 S(\mathbb{Q}_{\theta}, \mathbb{Q})|_{\mathbb{Q}=\mathbb{Q}_{\theta}}$ is called the scale function of the scoring rule *S*.

H.K. Olafsdottir, H. Rootzén and D. Bolin

Definition 2.1 (*Bolin & Wallin, 2023*). Let *S* be a proper scoring rule with respect to some class of probability measures \mathcal{P} on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, and assume that \mathcal{Q}_0 is a set of probability measures such that $\mathcal{Q} \subseteq \mathcal{P} \cap \mathcal{P}_S$. If $s(\mathbb{Q}_{\theta})$ exists for all $\mathbb{Q} \in \mathcal{Q}_0$, $\theta = (\mu, \sigma) \in \mathbb{R} \times \mathbb{R}^+$, and satisfies $s(\mathbb{Q}_{\theta}) \equiv \frac{1}{\sigma^2} s(\mathbb{Q}_{(0,1)})$ we say that *S* is locally scale invariant on \mathcal{Q} .

Bolin and Wallin (2023) showed that the CRPS has scale function $s(\mathbb{Q}_{\theta}) = \sigma s(\mathbb{Q}_{(0,1)})$ for $\theta = (\mu, \sigma)$, which means that the difference in Eq. (4) scales linearly with σ . The log-score on the other hand has a scale function $s(\mathbb{Q}_{\theta}) = \frac{1}{\sigma^2} H_{\mathbb{Q}}$ for a matrix $H_{\mathbb{Q}}$ that is independent of θ . The factor σ^2 therefore cancels in Eq. (4) so that the expression is independent of the scale.

A scoring rule that can be written as

$$S_g^{ker}(\mathbb{P}, y) := \frac{1}{2} E_{\mathbb{P}, \mathbb{P}}[g(X, X')] - E_{\mathbb{P}}[g(X, y)],$$
(5)

where g is a non-negative continuous negative definite kernel, is called a kernel score (Dawid, 2007). In order to construct locally scale invariant scoring rules, Bolin and Wallin (2023) introduced a generalisation of this class of scoring rules, the generalised kernel scores, which are proper scoring rules defined as

$$S_{g}^{h}(\mathbb{P}, y) := h(E_{\mathbb{P},\mathbb{P}}[g(X, X')]) + 2h'(E_{\mathbb{P},\mathbb{P}}[g(X, X')])$$
$$\times (E_{\mathbb{P}}[g(X, y)] - E_{\mathbb{P},\mathbb{P}}[g(X, X')]),$$
(6)

where *h* is any monotonically decreasing, convex, and differentiable function on \mathbb{R}^+ . For the particular choice $h(x) = -\frac{1}{2}\log(x)$ and g(x, y) = |x - y| in Eq. (6), one obtains the scoring rule *SCRPS*(\mathbb{P} , *y*) + 1, where the SCRPS is the scaled CRPS scoring rule

$$SCRPS(\mathbb{P}, y) := -\frac{E_{\mathbb{P}}[|X - y|]}{E_{\mathbb{P}, \mathbb{P}}[|X - X'|]} - \frac{1}{2}\log(E_{\mathbb{P}, \mathbb{P}}[|X - X'|]), \qquad (7)$$

which was shown to be locally scale invariant by Bolin and Wallin (2023). Another way of deriving this scoring rule is to use the fact that for any negative proper scoring rule $S(\mathbb{P}, x)$ on a set of probability measures \mathcal{P} , the transformed score

$$S_{S}^{trans}(\mathbb{P}, y) \coloneqq \frac{S(\mathbb{P}, y)}{|S(\mathbb{P}, \mathbb{P})|} - \log(|S(\mathbb{P}, \mathbb{P})|)$$
(8)

is also a proper scoring rule on \mathcal{P} (Bolin & Wallin, 2023).

3. Scaled weighted scoring rules

In this section we introduce the concept of local weight- and tail-scale invariance as less restrictive alternatives to local scale invariance. We then survey and discuss the use of the scaled wCRPS to score extremes and derive its scaling properties. In the final subsection, we propose alternative combined scoring rules for extremes, concentrating on a scoring rule based on the logarithmic score.

3.1. Local weight-scale invariance

When scoring extremes, we are mostly interested in the tail properties of the distributions, and in other cases, the interest might be on some specific range of values that is not necessarily extreme. Therefore, it might not be all that important to have full local scale invariance as long as this property holds in the region of interest. To make this precise, recall that Q is a family of locationscale transformations of probability measures in Q_0 . Let w be some weight function representing the region of interest as $\{x : w(x) > 0\}$, such that $P_{\mathbb{Q}}(w(X) > 0) > 0$ for all $\mathbb{Q} \in \mathcal{Q}_0$, and let \mathbb{Q}^w be the conditional distribution of $X^w = X|w(X) > 0$ if $X \sim \mathbb{Q}$. If X has density $f(\cdot)$, then X^w has density $f_w(\cdot) = f(\cdot)/P_{\mathbb{Q}}(w(X) > 0)$. A location-scale transformation $(a + bX)^w$ of X^w has density $f_w((\cdot - a)/b)/b$, which means that if $X^w \in \mathcal{Q}_0$, then we have $(a + bX)^w \in Q$. In other words, the conditional distributions of location-scale transformations are also location-scale transformations. Let Q^w denote the set of all these conditional distributions corresponding to $\mathbb{Q} \in \mathcal{Q}$. We assume for simplicity that the conditional distributions of the measures in Q_0 are in Q_0 , so that $\mathcal{Q}^w \subset \mathcal{Q}.$

As we are interested in the region $\{x : w(x) > 0\}$, we define the concept of weight-scale invariance by considering the scaling properties of *S* when restricted to Q^w .

Definition 3.1. Let *S* be a proper scoring rule with respect to some class of probability measures \mathcal{P} on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, and assume that \mathcal{Q}_0 is a set of probability measures such that $\mathcal{Q} \subseteq \mathcal{P} \cap \mathcal{P}_S$. Suppose that *w* is a weight function, and let $\mathcal{Q}^w \subset \mathcal{Q}$ be the set of conditional probability measures, as defined above. We say that *S* is a locally weight-scale invariant scoring rule with respect to a *w* on \mathcal{Q} if *S* is locally scale invariant on \mathcal{Q}^w .

Remark 3.2. Clearly, if *S* is locally scale invariant, it is also locally weight-scale invariant. However, as shown below, there are scoring rules that are locally weight-scale invariant but not locally scale invariant.

Recall that the scoring rule *S* is locally scale invariant if its scale function satisfies $s(\mathbb{Q}_{\theta}) = \frac{1}{\sigma^2} s(Q_{(0,1)})$. Considering the restriction of *S* to \mathcal{Q}^w , we have that *S* is locally weight-scale invariant if $s(\mathbb{Q}_{\theta}^w) = \sigma^{-2} s(\mathbb{Q}_{(0,1)}^w)$.

Remark 3.3. If the interest is focused on the region $w(\cdot) > 0$, we could simply use *S* restricted to Q^w as a scoring rule, which means that we only consider the score based on the conditional distributions. However, this disregards the probability of $w(\cdot) > 0$, which may be important in other situations. This probability is typically dependent on the scale, and local weight-scale invariance can thus be thought of as relaxation of local scale invariance where the scale dependence of the probability of $w(\cdot) > 0$ is ignored.

In the special case where the weight function is the indicator weight function w_u we refer to local weight-scale invariance as local tail-scale invariance. Specifically, we have the following definition.

H.K. Olafsdottir, H. Rootzén and D. Bolin

Definition 3.4. Let w_u be the indicator weight function in Eq. (3). Under the same assumptions as in Definition 3.1, we say that *S* is a locally tail-scale invariant scoring rule on Q if for each θ , there exists $u_{\theta} \in \text{supp}(\mathbb{Q}_{(0,1)}) \cap \text{supp}(\mathbb{Q}_{\theta})$, such that *S* is locally scale invariant on Q^{w_u} for every $u \ge u_{\theta}$. Similarly, we say that *S* is locally lower-tail-scale invariant on Q if *S* is locally scale invariant on Q^{1-w_u} for every $u \le u_{\theta}$.

3.2. Scaled weighted CRPS

For any kernel score we can use the construction in Eq. (6) for generalised kernel scores to construct a corresponding scaled version by using $h(x) = -\frac{1}{2}\log(x)$ as done for the SCRPS. By the following proposition, this can be done for the wCRPS.

Proposition 3.5. The wCRPS as defined in Eq. (2) is a kernel score.

This proposition was shown previously by Allen et al. (2022), and for completeness, a proof is given in Appendix C. Thus, by this proposition combined with the construction in Eq. (6), the scoring rule defined as

$$swCRPS(\mathbb{P}, y) := -\frac{E_{\mathbb{P}}[g_w(X, y)]}{E_{\mathbb{P}, \mathbb{P}}[g_w(X, X')]} - \frac{1}{2}\log(E_{\mathbb{P}, \mathbb{P}}[g_w(X, X')])$$
(9)

is proper, and can be a suitable alternative for evaluating extremes. As for the SCRPS, an alternative definition of the swCRPS would be to use Eq. (8) on the wCRPS. This definition was proposed by Vandeskog et al. (2022), and since the wCRPS is a kernel score, the two definitions in fact coincide, up to an additive constant.

No matter which of the arguments we use to define the swCRPS, it is not guaranteed that it is locally scale invariant. We now derive to what extent the swCRPS is scale invariant. The following assumption will be used on the distributions we consider.

Assumption 3.6. The Borel probability measure \mathbb{Q} on \mathbb{R} has density $\exp(\Psi)$ with respect to Lebesgue measure for some twice-differentiable function Ψ , such that the expectations $E_{\mathbb{Q}}[\Psi'(X)]$, $E_{\mathbb{Q}}[\Psi''(X)]$ and $E_{\mathbb{Q}}[(\Psi'(X))^2]$ are finite.

Under this assumption, we have the following result for the wCRPS and the swCRPS, which is proven in Appendix C.

Proposition 3.7. Let Q be a set of location-scale transformed probability measures \mathbb{Q}_{θ} satisfying Assumption 3.6. Then the following holds:

- (i) The wCRPS with indicator weight function w_u is neither locally scale invariant nor locally tail-scale invariant on Q.
- (ii) The swCRPS as defined in Eq. (9) with indicator weight function w_u is locally tail-scale invariant but not locally scale invariant on Q.

Other versions of the CRPS—the rCRPS and the rSCPRS, as introduced by Bolin and Wallin (2023)—are defined by inserting the kernel function

$$g_c(x, y) = \begin{cases} |x - y|, & |x - y| < c \\ 0, & \text{otherwise} \end{cases}$$

into the definition kernel scores in Eq. (5) and into Eq. (9). They showed that neither of these scores is locally scale invariant. However, one can show their local weight-scale invariance.

Proposition 3.8.

- (i) The rCRPS is not locally weight-scale invariant with respect to $w^{u}(y) = 1\{|y| < u\}$ for any u > 0.
- (ii) The rSCRPS is locally weight-scale invariant with respect to w^u for $u \leq \frac{c}{2}$.

The steps taken above by using the construction in Eq. (6) to create the swCRPS can be done on the more general threshold weighted kernel score proposed by Allen et al. (2022) and defined as

$$\mathsf{twS}_{\rho}(\mathbb{P}, y; v) = \frac{1}{2} E_{\mathbb{P}, \mathbb{P}}[\rho(v(X), v(X'))] - E_{\mathbb{P}}[\rho(v(X), v(y)] + \frac{1}{2}\rho(v(y), v(y)),$$

where ρ is a continuous negative-definite kernel, and $v(x)-v(x') = \int_{x'}^{x} w(t)dt$ is a measurable chaining function. This results in a scaled score

$$stwS_{\rho}(\mathbb{P}, y; v) = -\frac{E_{\mathbb{P}}[\rho(v(X), v(y)]}{E_{\mathbb{P},\mathbb{P}}[\rho(v(X), v(X'))]} - \frac{1}{2}\log(E_{\mathbb{P},\mathbb{P}}[\rho(v(X), v(X'))])$$

However, we focus on the more specific swCRPS to keep the exposition reasonably short, and leave further investigations of these scaled scores for future research.

As mentioned above, one main benefit of kernel scores is that the expected values can be computed with Monte Carlo methods without knowing the true distribution. However, under an assumed model distribution, the kernel scores can often be computed analytically. This is in particular the case for the generalised extreme value (GEV) distribution, as shown in Appendix A.

3.3. Local weight-scale invariance of the censored likelihood score

It follows from the definition that a sum of two proper scores is also proper, more generally that a weighted sum of several proper scores is proper, and additionally that if at least one of the summands is strictly proper, then the sum is also strictly proper. This can be used to propose alternative scores which may be less likely to suffer the problems of overweighting heavy tails, as discussed above for the quantile scores (see also, e.g., Lerch et al., 2017).

If one is interested in good predictions of values exceeding some threshold u, there are two complementary questions: (a) how well does a method predict the sizes of the excesses which occur, and (b) does the method give good prediction of the probability of exceedance? Here

H.K. Olafsdottir, H. Rootzén and D. Bolin

(a) can be approached by scoring only the values of the excesses which actually occur, using the conditional distribution of these, f(x)/(1 - F(u)). Question (b) is instead a binary prediction problem, and one can use any of the many scoring rules for such problems. Further, it seems reasonable to use related scoring rules for (a) and (b).

Diks et al. (2011) uses this idea to combine the logarithmic score for excesses:

$$S(\mathbb{P}, y) = 1\{y > u\} \log(f(y)/(1 - F(u)),$$

where *f* and *F* respectively denote the density and cumulative distribution functions corresponding to \mathbb{P} , with the binary logarithmic scoring rule for the exceedance of the threshold *u*,

$$S(\mathbb{P}, y) = 1\{y \le u\} \log(F(u)) + 1\{y > u\} \log(1 - F(u)),$$

to obtain the so-called censored likelihood score:

$$LS_u(\mathbb{P}, y) = 1\{y \le u\} \log(F(u)) + 1\{y > u\} \log(f(y)).$$
(10)

This is a proper scoring rule and is studied in simulations and application in the next two sections. We also have the following result, proven in Appendix C.

Proposition 3.9. Let Q be a set of location-scale transformed probability measures \mathbb{Q}_{θ} satisfying Assumption 3.6. Then the censored likelihood score in Eq. (10) is locally tail-scale invariant on Q, but in general not locally scale invariant.

The above type of weighted scoring rule was generalised in Holzmann and Klar (2017) as follows. First, let p_w represent the re-normalised density of p with respect to weight function w, i.e.

$$p_w(y) = \frac{w(y)p(y)}{\int_{\mathbb{R}} w(t)p(t)dt}$$

~

If S_0 is a proper scoring rule, then

$$owS(\mathbb{P}, y; w) = w(y)S_0(\mathbb{P}_w, y), \tag{11}$$

where \mathbb{P}_w is the distribution with density p_w , is a proper scoring rule. This is called an outcome weighted scoring rule by Allen et al. (2022).

Next, for a strictly proper scoring rule $\mathbf{s}(\alpha, z)$ for the success probability $\alpha \in (0, 1)$ of a binary outcome variable $z \in \{0, 1\}$, the score

$$S_{\mathbf{s}}(\mathbb{P}, \mathbf{y}; w) = w(\mathbf{y})\mathbf{s}\left(\int_{\mathbb{R}} p(t)w(t)dt, 1\right) + (1 - w(\mathbf{y}))\mathbf{s}\left(\int_{\mathbb{R}} p(t)w(t)dt, 0\right)$$
(12)

is said to be a localising proper weighted scoring rule for the density forecast *p*. Finally, adding Eqs. (11) and (12) yields a proper scoring rule:

$$S(\mathbb{P}, y, w) = S_{s}(\mathbb{P}, y; w) + \text{ow}S(\mathbb{P}, y; w).$$
(13)

With S_0 and **s** as log scores, the score \widehat{S} as defined in Eq. (13) becomes a censored likelihood with a general weight function w instead of the weight function w_u defined in Eq. (3).

International Journal of Forecasting xxx (xxxx) xxx

 Table 1

 Forecasting models for benchmarking, taken from Tail

lardat	et	al.	(2022).	The	true	model	is	Y Z	$\stackrel{d}{=}$	$\operatorname{Exp}(Z)$,
whoro	7	$\frac{d}{d}$		1	(1-ع	1 . 6 .	. (h		

$\begin{array}{lll} \hline Forecast & Density \\ \hline Ideal, \mathbb{P}_{ideal} & f_{Exp(Z)} \\ Climatological, \mathbb{P}_{clim} & f_{GP(1,\xi)} \\ \tau-Informed, \mathbb{P}_{\tau} & \tau_{fExp(Z)} + (1-\tau)f_{GP(1,\xi)} \\ Extremist, \mathbb{P}_{extr} & f_{Exp(Z/\nu)}, \nu > 1 \end{array}$	where $Z = Gainina(\xi)$	$, \xi , 1 > \xi > 0.$
$ \begin{array}{llllllllllllllllllllllllllllllllllll$	Forecast	Density
	Ideal, \mathbb{P}_{ideal} Climatological, \mathbb{P}_{clim} τ -Informed, \mathbb{P}_{τ} Extremist, \mathbb{P}_{extr}	$ \begin{array}{l} f_{Exp(Z)} \\ f_{GP(1,\xi)} \\ \tau f_{Exp(Z)} + (1-\tau) f_{GP(1,\xi)} \\ f_{Exp(Z/\nu)}, \ \nu > 1 \end{array} $

Proposition 3.10. Given that $\lambda(x : w(x) = 1) > 0$, where λ is the Lebesgue measure, the general censored likelihood score

$$CLog(\mathbb{P}, y; w) = w(y) \log \left(\int_{\mathbb{R}} p(t)w(t)dt \right)$$

+ $(1 - w(y)) \log \left(1 - \int_{\mathbb{R}} p(t)w(t)dt \right)$
+ $w(y) \log(p_w(y))$

is locally weight-scale invariant with respect to the indicator weight $\widetilde{w}(x) = 1\{w(x) = 1\}$.

Proposition 3.11. Suppose that S_0 and S_1 are locally weight-scale invariant with respect to w. Then:

- (i) If w is an indicator weight function, then owS as defined in Eq. (11) is locally weight-scale invariant with respect to w.
- (ii) $S_0 + S_1$ is locally weight-scale invariant with respect to w.

4. Simulation studies

This section contains two examples with simulated data, which illustrate the effect of scale dependence for extremes using the wCRPS and the swCRPS.

4.1. Benchmark example

We start by comparing the predictive ability of the CRPS against the SCRPS in terms of tail regimes. A comparison using the CRPS was previously done using a benchmark example in Taillardat et al. (2022). They considered the hierarchical model

$$\begin{cases} Z \stackrel{d}{=} \text{Gamma}(\xi^{-1}, \xi^{-1}) \\ Y|Z \stackrel{d}{=} \text{Exp}(Z) \end{cases}$$

where $1 > \xi > 0$, Gamma(α , β) denotes a gamma distribution with shape parameter α and rate parameter β , and Exp(δ) denotes an exponential distribution with scale parameter δ . Using this as the true model, four forecasting models are compared. The ideal model is an exponential model using observed values of *Z*. The extremist model instead underestimates the scale parameter of the exponential distribution, assuming it to be Z/ν , where $\nu > 1$. The climatological model has no information about *Z* but instead uses the unconditional and distributionally equivalent generalised Pareto distribution, GP(1, ξ), with scale parameter one and shape parameter ξ . And finally, the

H.K. Olafsdottir, H. Rootzén and D. Bolin

Table 2

Ratio of the mean CPRS and mean SCRPS with respect to the corresponding means for the ideal forecast for two different choices of ξ . The mean was taken over 10⁶ independently simulated observations.

Forecast	$\xi = 0.25$		$\xi = 0.50$	$\xi = 0.50$		
	CRPS	SCRPS	CRPS	SCRPS		
Ideal, P _{ideal}	100%	100%	100%	100%		
Extremist, $v = 1.1$	100.48%	100.41%	100.47%	100.39%		
0.75-Informed	100.89%	101.28%	102.14%	104.26%		
0.5-Informed	103.56%	103.76%	108.47%	109.93%		
Extremist, $v = 1.4$	106.67%	104.62%	106.64%	104.35%		
0.25-Informed	108.02%	107.20%	119.00%	116.31%		
Climatological, $GP(1, \xi)$	114.27%	113.67%	133.72%	131.58%		
Extremist, $v = 1.8$	122.87%	112.69%	122.83%	111.94%		



Fig. 1. Benchmark model simulation showing power of a two-sided pairwise t-test when scoring 1000 independently simulated stationary time series of length 1000 from the ideal model $Y|Z \sim \text{Exp}(Z), Z \sim \text{Gamma}(\xi^{-1}, \xi^{-1})$ using the extremist model with parameter $\nu \in (1.001, 2)$ and shape parameters $\xi = 0.25$ (left) and $\xi = 0.5$ (right).

 τ -informed model is a mixture distribution between the ideal and climatological model distributions. The models are listed in Table 1.

The CRPS for the extreme forecast and the τ -informed forecast are

$$CRPS(\mathbb{P}_{extr}, y) = -y - \frac{2\nu}{\delta} \exp\left(-\frac{\delta y}{\nu}\right) + \frac{3\nu}{2\delta}$$
(14)

and

$$CRPS(\mathbb{P}_{\tau}, y) = -y - \frac{\tau^2}{2\delta} - \frac{2\tau}{\delta} [exp(-\delta y) - 1] \\ + \frac{2(1-\tau)}{1-\xi} \left[1 - (1+\xi y)^{\frac{\xi-1}{\xi}} \right] \\ - \frac{(1-\tau)^2}{2-\xi} - \frac{2\tau(1-\tau)exp\left(\frac{\delta}{\xi}\right)}{\xi^{\frac{1}{\xi}}\delta^{\frac{\xi-1}{\xi}}}$$
(15)
$$\times \Gamma_u\left(\frac{\xi-1}{\xi}, \frac{\delta}{\xi}\right),$$

respectively, as shown by Taillardat et al. (2022). Here, $\Gamma_u(a, \tau) = \int_{\tau}^{\infty} t^{a-1} e^{-t} dt$ denotes the upper incomplete gamma function. From the CRPS of Eq. (14) and (15), one can compute the SCRPS by first noting that

$$E_{\mathbb{P}_{\text{extr}},\mathbb{P}_{\text{extr}}}[|X - X'|] = \frac{\tau}{\delta},$$

$$E_{\mathbb{P}_{\tau},\mathbb{P}_{\tau}}[|X - X'|] = 2\left(\frac{\tau}{\delta} - \frac{\tau^2}{2\delta} + \frac{1 - \tau}{1 - \xi} - \frac{(1 - \tau)^2}{2 - \xi}\right)$$

$$-\frac{2\tau(1-\tau)\exp\left(\frac{\delta}{\xi}\right)}{\xi^{\frac{1}{\xi}}\delta^{\frac{\xi-1}{\xi}}}\Gamma_u\left(\frac{\xi-1}{\xi},\frac{\delta}{\xi}\right)\right),$$

and $E_{\mathbb{P}_{\text{clim}},\mathbb{P}_{\text{clim}}}[|X - X'|] = \frac{2}{(2-\xi)(1-\xi)}$, and then using Eq. (7) to compute the SCRPS.

We simulate 10⁶ observations of the model for two choices of ξ : $\xi = 0.25$ and $\xi = 0.5$. The higher ξ is, the lower the variability in Z and thus the scale of Y|Z. Comparing the ratio of the mean scores of the different models to the corresponding means for the ideal model for the CRPS and SCRPS creates an ordering of the models, as seen in Table 2. Both scores order as expected within the τ -informed models and within the extremist models, with decreasing τ and increasing ν , respectively. However, the ordering within the two models depends on the choice of scoring rule. Hence, by choosing the scoring rule, one makes an implicit choice regarding which aspects of the model are important. For example, for $\xi = 0.25$, the climatological forecast is worse than the extremist forecast with v = 1.8 according to the CRPS but better according to the SCRPS, probably due to how the CRPS penalises the forecasts with larger uncertainty more than the SCPRS does, while the SCRPS prefers the models with the same relative error over the climatological model.

The sensitivity and power of the respective scores are also of interest for different values of ν . For a score *S*, the extreme model is compared against the ideal model on

H.K. Olafsdottir, H. Rootzén and D. Bolin



Fig. 2. Simulated mean (top) and standard deviation (bottom) of the score difference $S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta}) - S(\mathbb{P}, \mathbb{Q}_{\theta})$ using the wCRPS (left) and swCRPS (right) for two predictions of a random variable $X \sim \mathbb{Q}_{\theta} = \text{GEV}(\mu = 0, \sigma, \gamma = 0.12)$, with $\mathbb{P} = \text{GEV}(\mu = 0, 1.5\sigma, \gamma = 0.12)$, as functions of the scale parameter σ for different thresholds q(p), chosen as the *p*th quantile from \mathbb{Q} . Threshold q(0) results in the unweighted scores, CRPS and SCRPS.

a simulated dataset with 1000 observations using a twosided pairwise t-test. This is repeated 1000 times with different simulated datasets. The SCRPS better identifies the ideal model from the extremist model, as seen in Fig. 1, with higher power when comparing the ideal and the extremist model. Moreover, the power of the SCRPS does not change for increased values of ξ , whereas the power of the CRPS decreases.

4.2. Score dependence on scale and threshold

For an observation from $X \sim \mathbb{Q}_{\theta} = \text{GEV}(\mu = 0, \sigma, \gamma = 0.12)$, let us consider the effect of choosing a wrong model by computing the expected score difference $S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta}) - S(\mathbb{P}, \mathbb{Q}_{\theta})$ where $\mathbb{P} = \text{GEV}(\mu = 0, 1.5\sigma, \gamma = 0.12)$. Here, the value $\gamma = 0.12$ of the shape parameter is chosen to mimic the behaviour of the extreme rainfall events studied in Section 5. The mean and standard deviation of the score difference are computed from 50,000 simulated observations of *X*. For the weighted scores, the thresholds q(p)are obtained as the *p*th quantiles of \mathbb{Q}_{θ} , that is $q(p) = \mu + \frac{\sigma}{\gamma} \left((-\log p)^{-\gamma} - 1 \right)$ for p > 0, and q(0) represents the unweighted scores. Fig. 2 shows that for the wCRPS, differences increase with increased σ , whereas they remain stable for the swCRPS. The behaviour is similar for each choice of threshold. Variability increases with scale for the wCRPS but it remains constant for the swCRPS. Thus the local tail-scale invariance affects not only the mean score but also the score variance. The fact that the mean and variance of scores remain the same for different scales simplifies understanding and use of empirical distributions of scores from predictions with different scales. It can also help scoring scaled locations simultaneously, as seen in the following section.

4.3. Scaling effect on expected scores

As mentioned above, one often evaluates forecasts at multiple locations though average scores. If the individual scores are proper, their average is also a proper score. However, if the distributions at these locations have different scales, the use of a scale-dependent scoring rule leads to an implicit ranking of the importance of the different locations, in the sense that some locations might be more important to predict well to get a high average score.

As an example, consider two observations following GEV distributions that differ only in scale, $X_i \sim \mathbb{Q}_{\theta_i} = \text{GEV}(0, \sigma_i, 0.12), i = 1, 2$, with $\sigma_1 = 1.5$ and $\sigma_2 = 3$.

H.K. Olafsdottir, H. Rootzén and D. Bolin



Fig. 3. Simulated expected score $S(\mathbb{P}, \mathbb{Q})$ using the wCRPS and swCRPS for a pair (X_1, X_2) of random variables with $X_i \sim \text{GEV}(\mu_i, \sigma_i, \gamma)$ as functions of k_i , i = 1, 2, using a prediction that has the correct location parameters, μ_i , and scale parameters $\hat{\sigma}_i = k_i \sigma_i$, i = 1, 2. For the true model, $\mu_1 = \mu_2 = 0$, $\gamma = 0.12$, $\sigma_1 = 1.5$, and $\sigma_2 = 3$. The weight function was chosen as the 0.90 quantile for each score.

Using the predictions $\mathbb{P}_i = \text{GEV}(0, \hat{\sigma}_i, 0.12)$, where $\hat{\sigma}_i$ is an estimated scale parameter, the observations are scored using the scoring rule $S(\mathbb{P}, y) = \frac{1}{2} (S_1(\mathbb{P}_1, y_1) + S_2(\mathbb{P}_2, y_2))$, where S_1 and S_2 are either wCRPSs, or else swCRPSs with indicator weight functions w_{q_1} and w_{q_2} for the two observations. Here, q_1 and q_2 are chosen as the 90th quantiles of \mathbb{Q}_{θ_1} and \mathbb{Q}_{θ_2} , respectively. The expected score was computed as the mean of 100,000 simulations.

Fig. 3 shows that the expected average score $S(\mathbb{P}, \mathbb{Q})$ with $\hat{\sigma}_i = k_i \sigma_i$ for the wCRPS is more sensitive to changes in k_2 than in k_1 , meaning that it is more important to have a good prediction for X_2 . For the swCRPS, the scores are on the other hand quite symmetrical in k_1 , k_2 , so the predictions \mathbb{P}_1 , \mathbb{P}_2 are scored more equally. Fig. 3 further shows that the swCRPS is close to being locally scale invariant, whereas the wCRPS is not. Finally, note that the symmetry in k_1 , k_2 for all considered values indicates that the score difference does not depend on the scale even for large model misspecifications, which is not guaranteed to hold by the definition of local scale invariance. This is not uncommon, and for example also holds for Gaussian distributions, as can be seen in Bolin and Wallin (2023, Figure 1).

5. Case studies

This section uses data on water levels, precipitation, and air pollution to compare scoring rules in practice. The models used were fitted by maximising the log likelihood in the first two case studies and through the INLA approach (Rue et al., 2009) in the final study.

5.1. Extreme water levels

In this section, we consider data containing annual maximal water levels at five representative stations in the Great Lakes system: Lake St. Clair, Lake Michigan-Huron, Lake Ontario, Lake Superior, and Lake Erie. The dataset is provided by NOAA and dates from 1918 to 2020. With different sizes and depths of the lakes, the data differ in scale, making it interesting to see the effect of misspecifications when using the wCRPS and swCRPS.

First, a stationary GEV model was fitted to the five representative stations, assuming that the behaviour of the lake-wide average water levels has not changed during the observed time series. The estimated parameters are listed in Table 3. The shape parameter at all locations is negative, suggesting that the annual maxima follow a Weibull distribution. The density functions of the estimated GEV distributions are shown in Fig. D.10 in the appendix. Next, the $PGEV_{\lambda}$ model, with a trend in λ and temperature as covariate, was fitted to each of the stations. For further information on the PGEV model, see Olafsdottir et al. (2021). The temperature used was a LOWESS smoothing of the yearly average Northern Hemispheric temperature, obtained from NOAA (NOAA National Centers for Environmental Information, 2019). The fit suggested that Lake Michigan-Huron and Lake Superior did not have a trend in the expected number of extreme water levels, λ , but trends existed at the other three stations, Lake St. Clair, Lake Ontario, and Lake Erie, using significance level $\alpha = 0.05$. From the Lake System profile, the stationarity might be explained by flow from Lake Michigan, Lake Huron, and Lake Superior to the other lakes. The number of dams in the Great Lakes exceeds 7000, so the effect of them on the data is hard to visualise.

NOAA, Great Lakes Environmental Research Laboratory (2015) noted that "since September 2014, all of the Great Lakes have been above their monthly average levels for the first time since the late 1990s". However, they came to the same conclusion, that Lake Superior, Lake Michigan, and Lake Huron will probably remain stationary around the mean levels, while there might be non-stationarity in the remaining lakes in the Great Lakes system.

The above parameter estimates are used through simulation to compare the effect of parameter misspecification on the mean and standard deviation of the score differences. Further, since the stationary stations (Lake Superior and Lake Michigan–Huron) are also the ones that differ most in scale, these are used in the simulation study to compare how often one misspecified model is preferred over another equally misspecified model using different scores.

H.K. Olafsdottir, H. Rootzén and D. Bolin

International Journal of Forecasting xxx (xxxx) xxx

Table 3

Parameter estimates (standard deviations) using stationary GEV distribution.

Station	μ	σ	γ	Station id
St. Clair	175.108 (0.038)	0.349 (0.027)	-0.285 (0.065)	1
Michigan–Huron	176.469 (0.044)	0.395 (0.033)	-0.283 (0.082)	2
Ontario	74.990 (0.034)	0.322 (0.024)	-0.285 (0.053)	3
Superior	183.524 (0.019)	0.175 (0.014)	-0.404 (0.063)	4
Erie	174.280 (0.038)	0.355 (0.027)	-0.348 (0.060)	5



Fig. 4. Mean and standard deviations of score differences, Δ_i , at stations $i \in \{1, 2, 3, 4, 5\}$ with k = 1.5, using different types of CRPSs on simulated data using the estimated parameters from the stations listed in Table 3.

5.1.1. Simulation

Using the estimated parameters from Table 3, we simulate observations from the five different stations. For each station, 1000 independent stationary time series of length 100 are simulated. These time series have been scored for two models. The first model uses the true parameters μ , σ , and γ from the simulation, and the second model uses the true location and shape parameters μ and γ , but perturbs the scale parameter by a factor k = 1.5. The scale parameter of the second model therefore becomes 1.5σ . We see that for the parameters from station 4 (Lake Superior), both the mean score difference and its variation are smaller than for the other stations when using the wCRPS (Fig. 4). This difference between stations disappears when instead using the swCRPS on the same data. For a given perturbation factor k, the score difference at station i is defined as $\Delta_i := S_{\sigma_i} - S_{k\sigma_i}$ where $S_{\hat{\sigma}}$ is the score using estimated scale parameter $\hat{\sigma}$, and σ_i is the true scale parameter at station i (i.e., the parameter that was used when simulating the data). The factor k describes the estimated scale parameter's error relative to the true scale parameter. For scale invariant scores, the expected difference will only depend on the perturbation factor and not on the scale itself.

Assume that we want to compare two model predictions, each for two locations: Lake Michigan–Huron with $\sigma_2 = 0.395$, and Lake Superior with $\sigma_4 = 0.175$. Denote the models A and B, where

Model A: $\hat{\sigma}_2 = \sigma_2$, $\hat{\sigma}_4 = k\sigma_4$ Model B: $\tilde{\sigma}_2 = k\sigma_2$, $\tilde{\sigma}_4 = \sigma_4$

and k > 0 is a perturbation factor.



Fig. 5. Proportion of times model A was preferred over model B when simulating time series of length 100 when using the wCRPS and swCRPS, with individual shape parameter (left) and joint shape parameter (right). The score threshold is chosen as the p% quantile.



Fig. 6. Estimated scale parameter from a PGEV model with trend in the frequency parameter λ for stations in the northeastern U.S.

We want to compare the score $S_A := S_{\hat{\sigma}_2} + S_{\hat{\sigma}_4}$ against $S_B := S_{\widetilde{\sigma}_2} + S_{\widetilde{\sigma}_4}$. However, $S_A - S_B = \Delta_2 - \Delta_4$, so it suffices to compare the differences Δ_i at the stations. For this, 1000 time series of length 100 were simulated using the estimated parameters of Lake Superior and Lake Michigan–Huron. The perturbation factor was fixed at k =1.5. Fig. 5 shows how often model A was preferred over model B using different scores. Since the models both have one correct scale parameter and the same proportional error in the other parameter, models A and B should be chosen with equal probability if not influenced by scale. This corresponds to choosing A over B 50% of the time. Using the wCRPS gives a high bias toward model A, while the swCRPS yields a more fair comparison of models A and B, as expected. The deviation from the 0.5 line for large thresholds for the swCRPS is due to the difference in shape parameters at the simulated stations. If we assume that the shape parameters are the same for the two stations, this deviation disappears.

5.2. Extreme rainfall events and climate change

In this section we use the scoring rules introduced above to compare five models of extreme rainfall in the northeastern United States. The dataset is a part of NOAA Atlas 14 Volume 10, part 3 (NOAA's National Weather Service, 2019) and contains the annual maximum rain from 685 stations in the northeast with time series ranging from approximately 1900 to 2014. Only stations with at least 60 years of data are used. We call this dataset NE-data. The estimated scale parameters differ between stations (Fig. 6), meaning that the stations with higher scale parameters will be given greater weight if a scaledependent scoring rule such as the CRPS or wCRPS is used.

5.2.1. Model and inference details

We compared four models of annual maximal rainfall: (i) Gumbel, a GEV distribution with shape parameter zero; (ii) GEV, a GEV distribution with shape parameter

<u>ARTICLE IN PRESS</u>

H.K. Olafsdottir, H. Rootzén and D. Bolin

International Journal of Forecasting xxx (xxxx) xxx

Table 4

Mean scores for four different models for NE-data, with means obtained as the mean of the mean scores at the individual stations. The scores where t-tests showed significance between the PGEV_{λ} and GEV_{μ} are shown in bold.

	LS	CRPS	SCRPS	LSq		wCRPS		swCRPS	
				90%	99%	90%	99%	90%	99%
Gumbel	-1.148	-0.483	-1.048	-0.427	-0.0903	-0.0865	-0.010178	-0.1556	-1.531
GEV	-1.129	-0.481	-0.966	-0.419	-0.0856	-0.0860	-0.010201	-0.0795	0.893
GEV_{μ}	-1.108	-0.473	-0.958	-0.415	-0.0850	-0.0857	-0.010199	-0.0644	0.906
$PGEV_{\lambda}$	-1.107	-0.472	-0.956	- 0.414	-0.0847	-0.0854	-0.010195	-0.0603	0.901

different from zero; (iii) GEV_{μ} , a GEV distribution with trend in the location parameter, i.e. $\mu(t) = \mu_0 + \mu_1 t$; and (iv) PGEV_{λ} , a PGEV distribution with trend in the frequency parameter, i.e. $\ln \lambda(t) = \lambda_0 + \lambda_1 t$. For both GEV_{μ} and PGEV_{λ} , the covariate *t* is a LOWESS smoothing of the yearly average Northern Hemispheric temperature, obtained from NOAA (NOAA National Centers for Environmental Information, 2019).

The parameters of the GEV and the PGEV models are estimated by maximising the log likelihood. Both models assume a single regional shape parameter and stationwise location and scale parameters. The estimated scale parameter varies over space, motivating the use of locally (tail-) scale invariant scores. The thresholds for the weighted scores are determined empirically from the time series as the *p*th quantile.

Besides overall mean scores, we also consider the distribution of the scores. For station *i*, with N_i observations $\boldsymbol{y} = (y_1, \dots, y_{N_i})$, the mean score is

$$S_i(\mathbb{P}, \boldsymbol{y}) = \frac{1}{N_i} \sum_{j=1}^{N_i} S(\mathbb{P}_{ij}, y_j).$$

where \mathbb{P}_{ij} is a predictive distribution for observation j at station i. Let $\Delta_i(\mathbb{P}, \mathbb{Q}) = S_i(\mathbb{P}, \mathbf{y}) - S_i(\mathbb{Q}, \mathbf{y})$ denote the station-wise score difference between two predictions \mathbb{P} and \mathbb{Q} at station i. A positive difference means that prediction \mathbb{P} scored better, and a negative difference means that prediction \mathbb{Q} scored better. The standard way to evaluate whether \mathbb{P} is a better model than \mathbb{Q} for time series is to perform a Diebold–Mariano test (Diebold & Mariano, 1995) on the difference of this kind, that takes into account correlation in the data. However, for the station-wise differences, the spatial correlation is low, and we therefore perform pairwise t-tests of equal predictive performance to compare two models.

The variances of the station-wise average scores are affected by the varying scale parameters in the data, but also by the lengths of the time series, which are slightly different for the different stations. However, since only time series of length greater than 60 years are used, the latter difference is small.

5.2.2. Results

As shown in Table 4, the PGEV_{λ} model gives the highest overall mean scores, while the Gumbel model has the lowest score, for all statistically significant scores. Fig. 7 shows *p*-values from two-sided t-tests, where one can note that all scoring rules reject the null hypothesis except for very high thresholds, where only the LS_{*q*} score rejects.

Instead of only considering average scores, one can, as suggested by Taillardat et al. (2022), also check for the existence of trends by permuting covariates (in this paper, temperature) and computing prediction scores for the model which uses the permuted covariates instead of the ordered ones. In the absence of trends, these predictions should be similar to the ones which use the ordered covariates. By plotting the ordered average station scores from the original station data $\mathcal{D}_1 = (t_i, y_i)_{i=1}^n$ against the ordered average scores obtained by using the permuted covariates $\mathcal{D}_2 = (t_{\pi(i)}, y_i)_{i=1}^n$, one can see whether the scores from \mathcal{D}_1 are larger, which would point at the existence of a trend, or whether they distribute evenly around the 45 degree line, in which case a trend is unlikely. For the PGEV model with trend in frequency, Fig. 8 suggests the existence of a trend regardless of which scoring rule is used. Here, the same conclusion is reached regardless of the score chosen. In particular, the rankings of the models do not change if we focus only on high values. Whether this is the case is application-dependent, and Section 5.3 shows an example where the conclusions differ when only focusing on high values.

5.3. Particle matter concentration

So far, the examples have focused on extremes, due to the natural interest in predictions above a certain threshold. However, there are many cases where there is a reason to focus on specific regions which are not necessarily extreme. For example, high concentrations of particle matter less than 10 nm in diameter (PM10) have negative health impact. According to the WHO (Geneva: World Health Association, 2021), daily PM10 should not exceed 45 μ g/m³. Moreover, based on studies on the connection between long-term PM10 and non-accidental mortality, the WHO found that the long-term level of PM10 should not exceed 15 μ g/m³. This is yet a recommendation, but it might be of interest to evaluate how well models can predict PM10, focusing on values above 15 μ g/m³. The assumption is that PM10 concentrations below 15 μ g/m³ have less health implications, and thus the prediction of lower values is not as important for the model evaluation. As an illustration, we focus on concentrations above three distinct thresholds $-15 \ \mu g/m^3$, 30 μ g/m³, and 45 μ g/m³–when comparing predictions of PM10 concentrations in the Piedmont region in Italy. We use data from Cameletti et al. (2013), containing observations for 182 days from October 2005 to March 2006 at 24 locations. The data are illustrated in Fig. 9.

H.K. Olafsdottir, H. Rootzén and D. Bolin

International Journal of Forecasting xxx (xxxx) xxx



Fig. 7. *P*-values from t-test station-wise score differences $\Delta_i(\mathbb{P}, \mathbb{Q})$ for predictions \mathbb{P} and \mathbb{Q} , where (A) $\mathbb{P} = GEV$, $\mathbb{Q} = GEV_{\gamma=0}$, (B) $\mathbb{P} = PGEV_{\lambda}$, $\mathbb{Q} = GEV_{\gamma=0}$, (C) $\mathbb{P} = PGEV_{\lambda}$, $\mathbb{Q} = GEV_{\lambda}$, $\mathbb{Q} = GEV_{\mu}$. The 0.05 level is marked with a red dashed line. Note that the *y*-scales are different in the different plots.



Fig. 8. Sorted average station scores for a PGEV model with trend in λ plotted for the original dataset D_1 against permuted dataset D_2 . The colour represents whether the points fall above (green) or below (orange) the diagonal. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Denoting the n = 4368 observations of log-PM10 by $y_i, i = 1, ..., n$, Cameletti et al. (2013) proposed using a model of the form

$$y_i = \sum_{m=1}^{M} W_m(s_i, t_i)\beta_m + u(s_i, t_i) + e_i.$$

Here, W_m denotes covariates, β_m regression coefficients, and e_i independent centred Gaussian variables representing measurement noise. Further, u(s, t) is a spatiotemporal Gaussian random field with a separable covariance function.

Since separability is rarely a realistic assumption for spatiotemporal statistical models, we compare the pre-



Fig. 9. Position of the 24 stations in the Piedmont region along with observations of PM10 concentration on March 13, 2006 (right), and a histogram showing the observations from the 24 stations at the 182 different dates (left).

Table	5
-------	---

Mean scores from one-step-ahead cross-validation used to compare PM10 predictions in the Piedmont data for separable and non-separable spacetime models, as in Krainski (2023). The higher score is represented in bold. The thresholds used for the weighted scores are 15 μ g/m³, 30 μ g/m³, and 45 μ g/m³.

	CRPS	SCRPS	wCRPS			swCRPS			
			$15 \ \mu g/m^3$	$30 \ \mu g/m^3$	$45 \ \mu g/m^3$	$15 \ \mu g/m^3$	$30 \ \mu g/m^3$	$45 \ \mu g/m^3$	
Separable Non-separable	-0.2243 - 0.2232	-0.5792 - 0.5750	-0.1987 - 0.1983	-0.1600 - 0.1590	-0.1151 - 0.1590	- 0.5516 -0.5863	- 0.4960 -1.1672	- 0.1452 -0.2162	

dictions based on this model with those of a model where u is replaced by a Gaussian random field with a nonseparable covariance function. Specifically, we use the critical diffusion model of Lindgren et al. (2023), which is a Gaussian random field defined as a solution to a stochastic partial-differential equation (SPDE). Both models are fitted to the data using the code provided by Krainski (2023).

To compare the models in terms of predictive performance, a one-step-ahead cross-validation was performed by extending the cross-validation function of the R package rSPDE (Bolin & Simas, 2023) to include the swCRPS and wCRPS, with the test set at fold *i* defined as all observations at time t_{i+1} , and the test set as all observations up to that time, for i = 1, ..., 181. The predictions are compared in terms of CRPS, SCRPS, wCRPS, and swCRPS. The results of the best scores are shown in Table 5. In line with the results found in Lindgren et al. (2023), the CRPS and SCRPS indicate that the non-separable model has better predictive performance than the separable model. However, for the weighted scores, the wCRPS indicates better performance of the non-separable model, while the swCRPS indicates better performance of the separable model. Thus, this application is an example of where the choice of scoring rule affects the rankings. In particular, we see that the swCRPS indicates that the separable model might be preferable when focusing on harmful levels of PM10.

6. Conclusion and discussion

Desirable properties such as straightforward computation of scores by Monte Carlo approximation have made the CRPS a popular alternative to the logarithmic score. However, using the CRPS requires sacrificing the local scale invariance of the logarithmic score. This can cause different observations, say at different spatial locations, not to be equally important for average scores. To account for this, a common choice in the literature is to use socalled skill scores $(S_n - S_n^{ref})/S_n^{ref}$ or $(S_n - S_n^{ref})/(S_n^{opt} - S_n^{ref})$, where S_n^{ref} and S_n^{opt} are the scores for a reference method and the hypothetical optimal score, respectively. However, these scores are often improper, even when they are based on a proper scoring rule (Gneiting & Raftery, 2007).

H.K. Olafsdottir, H. Rootzén and D. Bolin

Bolin and Wallin (2023) addressed this issue by introducing the SCRPS, a proper scoring rule which retains most of the desirable properties of the CRPS yet is also locally scale invariant. We extend this construction to the weighted CRPS to obtain the swCRPS and investigate the properties of this score. We show that neither the wCRPS nor the swCRPS is scale invariant, but that the swCRPS is locally tail-scale invariant, meaning that, conditioned on exceeding the threshold used in the weighting, the score

is locally scale invariant. Through simulation studies we showed that for a number of extreme value models, swCRPS works in a similar way as the SCRPS does for full datasets. In particular, the mean and variability of score differences remain similar when scale parameters are changed, and average scores over random variables with different scales keep the relative error in scale parameter estimation approximately equal. These properties can be important, for example when evaluating weather models that operate on different scales at different spatial locations, such as temperature, water levels, or rainfall amounts at different measuring stations.

As an example, we used different scores to evaluate five models on annual maximum rainfall data in the northeastern USA. The comparison led to similar conclusions for all scores, with the best tested model being a PGEV model with a temperature-wise trend in the frequency parameter, meaning that for this case study, we come to the same conclusion when only considering the tail of the distribution. This was not the case for the particle matter application, where focusing on high values changed the model ranking.

It should be noted that scale dependence can sometimes be preferred, as mentioned by Bolin and Wallin (2023). In such situations, the wCRPS would likely be a better choice than the swCRPS. However, we believe that such situations are rare for typical applications in extremes. Moreover, the LS_q seems to have the best power, as expected, but this score might not be possible to compute for more complicated models. In such cases the swCRPS can be a good alternative.

Issues regarding tail-weighted scores have been raised as the forecaster's dilemma (Lerch et al., 2017). The lack of information from observations below a certain threshold can unintentionally draw the score towards the wrong model based on the weight of the models' tail behaviour. This has been addressed by also including the unweighted CRPS by using a weight function of the form $a + bw_u(x)$ (Thorarinsdottir & Schuhen, 2018). Another solution might be simply not to choose an excessively high threshold u.

Weighted scores are not solely relevant for extremes but can be applied whenever predictions within a specific region are of higher importance, as seen for the PM10 concentration model, where a Gaussian spatiotemporal model was used to predict the concentration of PM10 in the Piedmont region of Italy (Cameletti et al., 2013; Lindgren et al., 2023). Moreover, the regions of interest can be defined through weight functions other than the indicator weight function $w_u(y) = 1\{y \ge u\}$. The notation of local weight-scale invariance extends the concept of local tail-scale invariance to accommodate for local scale invariance within these types of regions.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank Thordis Thorainsdottir and two anonymous reviewers for very helpful comments which have led to substantial improvements of the paper.

Appendix A. Extreme value models and scoring rules

When considering block maxima data, such as annual maxima of daily precipitation data, classical extreme value theory (see e.g. Coles, 2001) shows that the generalised extreme value (GEV) distribution is a suitable model for the data. Different variations of this approach exist, as shown below.

A.1. Extreme value theory models

The GEV distribution has location, scale, and shape parameters (μ , σ > 0, and γ respectively) and the distribution function

$$F_{\text{GEV}}(x) = \begin{cases} \exp\left(-\exp\left(-\frac{x-\mu}{\sigma}\right)\right) & \text{if } \gamma = 0, \\ \exp\left\{-\left(1+\frac{\gamma}{\sigma}(x-\mu)\right)^{-1/\gamma}\right\} & \text{if } \gamma \neq 0. \end{cases}$$

In this model, one can introduce trends in the parameters, for example to model climate change. Further, one can use a re-parameterisation, PGEV, of the GEV distribution in terms of parameters λ , σ_u , γ , u, with

$$\mu_{\text{PGEV}} = u + \frac{(\lambda^{\gamma} - 1)\sigma_u}{\gamma}, \quad \sigma_{\text{PGEV}} = \sigma_u \lambda^{\gamma}, \quad \gamma_{\text{PGEV}} = \gamma.$$

Here the parameter λ describes the frequency of exceedances of a high level *u*, and σ_{PGEV} is a scale parameter of the distribution of sizes of excesses of *u*. A reason for this re-parameterisation is that it gives a clearer physical understanding of the behaviour of extreme events. The distribution function of the PGEV can be written as

$$F_{\text{PGEV}}(x) = \exp\left\{-\lambda \left(1 + \frac{\gamma}{\sigma_u}(x-u)\right)^{-1/\gamma}\right\}.$$

For details and more information, see Olafsdottir et al. (2021). For models without trends in parameters, the PGEV distribution is the same as the GEV distribution. However, the trends in λ for the PGEV model and in μ for the GEV model behave differently.

When \mathbb{P} is the GEV or PGEV distribution, and $\gamma < 1$, the CRPS has a closed-form expression (see e.g. Friederichs & Thorarinsdottir, 2012),

H.K. Olafsdottir, H. Rootzén and D. Bolin

International Journal of Forecasting xxx (xxxx) xxx

 $CRPS(\mathbb{P}, y)$

$$= \begin{cases} -(y - \mu + \frac{\sigma}{\gamma})(2F(y) - 1) \\ + \frac{\sigma}{\gamma}(2^{\gamma}\Gamma(1 - \gamma) - 2\Gamma_{l}(1 - \gamma, -\ln F(y))) \\ \text{if } \gamma \neq 0, \\ (y - \mu) + 2\sigma Ei(\ln F(y)) - \sigma(C - \ln 2) \\ \text{if } \gamma = 0, \end{cases}$$

where $\Gamma(a) = \int_0^\infty t^{a-1} e^{-t} dt$ is the gamma function, $\Gamma_l(a, \tau) = \int_0^\tau t^{a-1} e^{-t} dt$ is the lower incomplete gamma function, $Ei(x) = \int_{-\infty}^x e^t / t dt$ is the exponential integral, and *C* is the Euler–Masceroni constant. For $\gamma \ge 1$, the CRPS does not exist. In Appendix B we derive the corresponding closed-form expressions for the wCRPS.

Through the following reformulation of Eq. (9),

$$swCRPS(\mathbb{P}, y) = \frac{wCRPS(\mathbb{P}, y)}{E_{\mathbb{P}, \mathbb{P}}[g_w(X, X')]} - \frac{1}{2}\ln(E_{\mathbb{P}, \mathbb{P}}[g_w(X, X')]) - \frac{1}{2},$$

we can also evaluate $swCRPS(\mathbb{P}, y)$ for the GEV distribution through the closed-form expression for the wCRPS in combination with the expression

$$E_{\mathbb{P},\mathbb{P}}[g_w(X,X')] = -2\left(q - \mu + \frac{\sigma}{\gamma}\right)F(q)(1 - F(q)) + 2\frac{\sigma}{\gamma}\left[2^{\gamma}\Gamma_l(1 - \gamma, -2\ln(F(q)))\right] - 2\Gamma_l(1 - \gamma, -\ln(F(q)))\right].$$

See Appendix B for the derivation of these expressions.

Appendix B. Closed-form expressions for GEV scores

Let \mathbb{P} be a distribution with CDF *F* and PDF *f*, and let w_q be the indicator weight function as described in Eq. (3). Then

$$E_{\mathbb{P}}[g_w(X, y)] = \int_{\mathbb{R}} g_w(x, y)f(x)dx$$

= $(1 - w_q(y))\int_q^{\infty} (x - q)f(x)dx$
+ $w_q(y)\left(\int_{-\infty}^q (y - q)f(x)dx + \int_q^{\infty} |x - y|f(x)dx\right)$
= $\int_q^{\infty} xf(x)dx - q(1 - F(q)) + w_q(y)$
 $\times \left(y(2F(y) - 1) - q(2F(q) - 1) - 2\int_q^y xf(x)dx\right),$

and

$$E_{\mathbb{PP}}[g_w(X, X')] = \int_{\mathbb{R}} \int_{\mathbb{R}} g_w(x, y) f(x) dx f(y) dy$$

$$= 4 \int_q^{\infty} xF(x) f(x) dx$$

$$- 2 \int_q^{\infty} xf(x) dx - 2qF(q)(1 - F(q)).$$
(B.1)

Using the formulation in Eq. (2), we have

$$\begin{split} \text{wCRPS}(\mathbb{P}, y) &= \frac{1}{2} E_{\mathbb{PP}} F[g_w(X, X')] - E_{\mathbb{P}}[g_w(X, y)] \\ &= 2 \int_q^\infty x F(x) f(x) dx - \int_q^\infty x f(x) dx - q F(q) (1 - F(q)) \\ &- \int_q^\infty x f(x) dx + q (1 - F(q)) \\ &- w_q(y) \left(y(2F(y) - 1) - q(2F(q) - 1) - 2 \int_q^y x f(x) dx \right) \\ &= -2 \int_{q \lor y}^\infty x f(x) dx + 2 \int_q^\infty x F(x) f(x) dx \\ &- (q \lor y) (2F(q \lor y)) - 1) + q F(q)^2. \end{split}$$
(B.2)

The expressions in Eq. (B.1) and (B.2) can be rewritten in terms of quantiles of the distribution *F*, since we have $\int_a^{\infty} xf(x)dx = \int_{F(a)}^1 F^{-1}(t)dt$ and $\int_a^{\infty} xF(x)f(x)dx = \int_{F(a)}^1 F^{-1}(t)tdt$. This gives

$$E_{\mathbb{PP}}[g_w(X, X')] = 4 \int_{F(q)}^{1} F^{-1}(t) t dt$$

- $2 \int_{F(q)}^{1} F^{-1}(t) dt - 2qF(q)(1 - F(q)).$

and

wCRPS(
$$\mathbb{P}, y$$
) = $-2 \int_{F(q \lor y)}^{1} F^{-1}(t) dt + 2 \int_{F(q)}^{1} F^{-1}(t) t dt - (q \lor y)(2F(q \lor y)) - 1) + qF(q)^2.$

Let us now use these results to evaluate the scoring rules for the GEV distribution. We first consider the case $\gamma = 0$, corresponding to the Gumbel distribution, and then the case $\gamma \neq 0$.

B.1. Scoring rules for the Gumbel distribution

For the case $\gamma = 0$, the GEV distribution has CDF $F(x) = \exp(-\exp(-(x - \mu)/\sigma))$ with inverse $F^{-1}(x) = \mu - \sigma \ln(-\ln x)$). Therefore,

$$\int_{F(q\vee y)}^{1} F^{-1}(t)dt = \mu(1 - F(q \vee y)) - \sigma \int_{F(q\vee y)}^{1} \ln(-\ln t)dt,$$

$$\int_{F(q)}^{1} F^{-1}(t)tdt = \frac{1}{2}\mu(1 - F(q)^{2}) - \sigma \int_{F(q)}^{1} t \ln(-\ln t)dt,$$

which gives

$$wCRPS(\mathbb{P}, y) = ((q \lor y) - \mu) - \sigma(C - \ln(2))$$

$$-\sigma\left(Ei\left(-2\exp\left(-\frac{q-\mu}{\sigma}\right)\right)\right)$$
$$-2Ei\left(-\exp\left(-\frac{q\vee y-\mu}{\sigma}\right)\right),$$

and

$$E_{\mathbb{PP}}[g_w(X, X')] = 2\sigma \ln(2) - 2(q - \mu)(1 - F(q))$$
$$- 2\sigma \left(Ei\left(-2\exp\left(\frac{q - \mu}{\sigma}\right)\right) \right)$$
$$- Ei\left(-\exp\left(\frac{q - \mu}{\sigma}\right)\right) \right).$$

H.K. Olafsdottir, H. Rootzén and D. Bolin

International Journal of Forecasting xxx (xxxx) xxx

B.2. Scoring rules for the GEV distribution with non-zero shape parameter

For the case $\gamma \neq 0$, the GEV distribution has CDF $F(x) = \exp\left[-\left(1 + \frac{\gamma}{\sigma}(x-\mu)\right)^{-1/\gamma}\right]$ with inverse $F^{-1}(x) = \mu - \frac{\sigma}{\gamma}\left(1 - (-\ln x)^{-\gamma}\right)$. Therefore,

wCRPS(F, y) =
$$\left(q - \mu + \frac{\sigma}{\gamma}\right)F(q)^2$$

- $\left((y \lor q) - \mu + \frac{\sigma}{\gamma}\right)(1 - 2F(y \lor q))$
- $\frac{\sigma}{\gamma}\left[2^{\gamma}\Gamma_l(1 - \gamma, -2\ln(F(q)))\right]$
- $2\Gamma_l(1 - \gamma, -\ln(F(y \lor q)))\right],$

and

$$E_{\mathbb{PP}}[g_w(X, X')] = -2\left(q - \mu + \frac{\sigma}{\gamma}\right)F(q)(1 - F(q))$$
$$+ 2\frac{\sigma}{\gamma}\left[2^{\gamma}\Gamma_l(1 - \gamma, -2\ln(F(q)))\right]$$
$$- 2\Gamma_l(1 - \gamma, -\ln(F(q)))],$$
where $\Gamma_l(q, \tau) = \int_{\tau}^{\tau} t^{a-1}e^{-t}dt$

where $I_{l}(a, \tau)$ Jou

Appendix C. Proofs

Proof of Proposition 3.5. Let $g_w(x, y) := |\int_y^x w(t)dt|$, where w(x) is a non-negative function. Then g_w : $\Omega \times$ $\Omega \to \mathbb{R}$ is a symmetric real-valued function. Moreover,

$$\sum_{i=1}^n \sum_{j=1}^n a_i a_j g_w(x_i, x_j) \le 0$$

for all n, all $a_1, \ldots, a_n \in \mathbb{R}$ s.t. $\sum_{i=1}^n a_i = 0$, and all $x_1, \ldots, x_n \in \Omega$, since g(x, y) = |x - y| is negative definite. Hence, the kernel g_w is negative definite and the weighted CRPS,

wCRPS(
$$\mathbb{P}, y$$
) = $\frac{1}{2} E_{\mathbb{P},\mathbb{P}}[g_w(X, X')] - E_{\mathbb{P}}[g_w(X, y)],$

is indeed a kernel score.

Proof of Proposition 3.7.

Note that for any \mathbb{Q} for which Assumption 3.6 holds, the assumption also holds for \mathbb{Q}^{w_u} , where $u \in \text{supp}(\mathbb{Q})$, and w_u is the indicator weight function.

We start by proving (i). Assume that \mathbb{Q}_{θ} is a probability measure satisfying Assumption Assumption 3.6. Let $S(\mathbb{Q}_{\theta}, y) = \frac{1}{2} E_{\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta}}[g_w(X, Y)] - E_{\mathbb{Q}_{\theta}}[g_w(X, y)]$ denote the wCRPS, where q_{θ} is the density of \mathbb{Q}_{θ} . By Taylor expansion around θ , we have that

$$S(\mathbb{Q}_{\theta+t\sigma\boldsymbol{r}},\mathbb{Q}_{\theta}) = S(\mathbb{Q}_{\theta},\mathbb{Q}_{\theta}) + t\sigma\boldsymbol{r}^{T}\nabla_{\theta}S(\mathbb{Q}_{\theta},\mathbb{Q}_{\theta}) + \frac{1}{2}t^{2}\sigma^{2}\boldsymbol{r}^{T}\nabla_{\theta}^{2}S(\mathbb{Q}_{\theta},\mathbb{Q}_{\theta})\boldsymbol{r} + o(t^{2}).$$
(C.1)

From Assumption 3.6, $s(\theta) = \nabla^2_{\theta} S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta})|_{\mathbb{Q}=\mathbb{Q}_{\theta}}$ exists and is continuous, and we have $\nabla_{\theta} S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta}) = 0$, since *S*

is proper. Hence,

$$S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta}) - S(\mathbb{Q}_{\theta+t\sigma\mathbf{r}}, \mathbb{Q}_{\theta})$$

= $-\frac{1}{2}t^{2}\sigma^{2}\mathbf{r}^{T}\nabla_{\theta}^{2}S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta})\mathbf{r} + o(t^{2})$
= $\frac{1}{4}t^{2}\sigma^{2}\mathbf{r}^{T}\nabla_{\theta}^{2}E_{\mathbb{Q}_{\theta},\mathbb{Q}_{\theta}}[g_{w}(X, Y)]\mathbf{r} + o(t^{2})$

We can now follow the steps in the proof of Lemma 1 in Bolin and Wallin (2023), where we replace their kernel $g_c(x, y)$ with our weighted kernel $g_w(x, y)$, since the only thing required in the proof is that the kernel is a positive negative-definite kernel and that $g_w(x, y) \leq g(x, y)$, both of which hold whenever $|w(t)| \leq 1$ for all $t \in \mathbb{R}$. This holds certainly for our indicator weight function w_a but even for other choices of w. This results in

$$\nabla_{\theta} E_{\mathbb{Q}_{\theta},\mathbb{Q}}[g_w(X,Y)] = -\sigma^{-1} E_{\mathbb{Q},\mathbb{Q}}[g_w(\sigma X + \mu, Y)v(X)], \quad (C.2)$$

and

$$\nabla_{\theta}^{2} E_{\mathbb{Q}_{\theta},\mathbb{Q}_{\theta}}[g_{w}(X,Y)]$$

= $\sigma^{-2} E_{\mathbb{Q},\mathbb{Q}}[g_{w}(\sigma X + \mu, \sigma Y + \mu)H(X,Y)],$ (C.3)

where v(X) is a vector and H(X, Y) is a 2 × 2 matrix, both independent of θ , and

$$g_w(\sigma x + \mu, \sigma y + \mu) = \left| \int_{\sigma y + \mu}^{\sigma x + \mu} w(t) dt \right|$$
$$= \sigma \left| \int_y^x w(\sigma \tau + \mu) d\tau \right|$$

For $w = w_q$, we have

$$g_{w}(\sigma x + \mu, \sigma y + \mu) = \sigma \left| \int_{y}^{x} 1\left\{ \tau \ge \frac{q - \mu}{\sigma} \right\} d\tau \right|$$
$$= \begin{cases} \sigma |x - y| & \text{if } x, y \ge \frac{q - \mu}{\sigma}, \\ \sigma |y - \frac{q - \mu}{\sigma}| & \text{if } y \ge \frac{q - \mu}{\sigma} \ge x, \\ \sigma |x - \frac{q - \mu}{\sigma}| & \text{if } x \ge \frac{q - \mu}{\sigma} \ge y, \\ 0 & \text{otherwise.} \end{cases}$$
(C.4)

Since we cannot let *w* depend on $\theta = (\mu, \sigma)$, we cannot choose any w such that

$$g_w(\sigma x + \mu, \sigma y + \mu) = \sigma g_w(x, y).$$

However, letting $k = \frac{q-\mu}{\sigma}$, we have that $g_w(\sigma x + \mu, \sigma y + \mu) = \sigma |x - y|$ for all $x, y \ge k$.

Then the wCRPS has scale function

 $s(\mathbb{Q}_{\theta}) = \sigma^{-1} E_{\mathbb{Q},\mathbb{Q}}[H_{\mathbb{Q}}(X, Y)h(X, Y)],$

where $H_{\mathbb{Q}}(X, Y)$ is a 2 × 2 matrix independent of θ , and h(X, Y) is a function dependent on θ . Furthermore, there exists a constant $k < \infty$ such that h(X, Y) is independent of θ for all X, Y > k. From this scale function we directly see that the wCRPS is neither locally scale invariant nor locally tail-scale invariant.

We now prove (ii). Let S be the swCRPS defined in Eq. (9). As in the proof of statement (i), we perform the Taylor expansion shown in Eq. (C.1). Since S is proper, $\nabla_{\theta} S(\mathbb{Q}_{\theta}, \mathbb{Q}_{\theta}) = 0$ and we only need to consider the term $\nabla^2_{\theta} S(\mathbb{Q}_{\theta}, \mathbb{Q})|_{\mathbb{Q}=\mathbb{Q}_{\theta}}$ that exists and is continuous.

For simplified notation, let $E_{P,Q} = E_{\mathbb{P}_{\theta},\mathbb{Q}}[g(X, Y)]$, $E_{\dot{P},Q} = \nabla_{\theta} E_{\mathbb{P}_{\theta},\mathbb{Q}}[g(X,Y)], E_{Q,\dot{P}} = \nabla_{\theta} E_{\mathbb{Q},\mathbb{P}_{\theta}}[g(X,Y)], \text{ and}$

H.K. Olafsdottir, H. Rootzén and D. Bolin

 $E_{P,Q} = \nabla_{\theta}^2 E_{\mathbb{P}_{\theta},Q}[g(X, Y)]$. One can show that

$$\begin{aligned} \nabla^2_{\theta} \log(E_{P,P}) &= -\frac{2E_{\dot{P},P}E_{\dot{P},P}^T}{E_{P,P}^2} + \frac{E_{\dot{P},\dot{P}}}{E_{P,P}} + \frac{E_{\ddot{P},P}}{E_{P,P}}, \\ \nabla^2_{\theta} \frac{E_{P,Q}}{E_{P,P}} &= -\frac{2E_{\dot{P},Q}E_{\dot{P},P}^T}{E_{P,P}^2} - \frac{2E_{P,Q}E_{\dot{P},P}^T}{E_{P,P}^2} - \frac{2E_{P,Q}E_{\dot{P},\dot{P}}^T}{E_{P,P}^2} \\ &+ \frac{2^3E_{\dot{P},P}E_{\dot{P},P}^T}{E_{P,P}^3} + \frac{E_{\ddot{P},P}}{E_{P,P}} - \frac{2E_{P,Q}E_{\dot{P},P}^T}{E_{P,P}^2}. \end{aligned}$$

Evaluating these together at Q = P yields

$$\nabla^2_{\theta} S(\mathbb{P}_{\theta}, \mathbb{Q})|_{\mathbb{Q}=\mathbb{P}_{\theta}} = \frac{1}{E_{P,P}} E_{\dot{P},\dot{P}} - \frac{2}{E_{P,P}^2} E_{\dot{P},P} E_{P,\dot{P}}^T$$

Inserting (C.2) and (C.3) instead of $E_{\dot{P},P}$ and $E_{\dot{P},\dot{P}}$ into the equation above results in

$$\begin{split} \mathbf{s}(\mathbb{Q}_{\boldsymbol{\theta}}) &= -\frac{1}{2} \nabla_{\boldsymbol{\theta}}^{2} S(\mathbb{Q}_{\boldsymbol{\theta}}, \mathbb{Q})|_{\mathbb{Q}} = \mathbb{Q}_{\boldsymbol{\theta}} \\ &= -\frac{\sigma^{-2}}{2} \left(\frac{E_{\mathbb{Q},\mathbb{Q}}[g_{w}(\sigma X + \mu, \sigma Y + \mu)H(X, Y)]}{E_{\mathbb{Q},\mathbb{Q}}[g_{w}(\sigma X + \mu, \sigma Y + \mu)]} \\ &- \frac{2E_{\mathbb{Q},\mathbb{Q}}[g_{w}(\sigma X + \mu, \sigma Y + \mu)v(X)]E_{\mathbb{Q},\mathbb{Q}}[g_{w}(\sigma X + \mu, \sigma Y + \mu)v(X)]^{T}}{E_{\mathbb{Q},\mathbb{Q}}[g_{w}(\sigma X + \mu, \sigma Y + \mu)]^{2}} \right). \end{split}$$

and for $u \ge k$,

$$\begin{split} s(\mathbb{Q}_{\theta}^{w_{u}}) &= -\frac{\sigma^{-2}}{2} \left(\frac{\sigma E_{\mathbb{Q}^{w_{u}},\mathbb{Q}^{w_{u}}}[|X-Y|H(X,Y)]}{\sigma E_{\mathbb{Q}^{w_{u}},\mathbb{Q}^{w_{u}}}[|X-Y|]} \\ &- \frac{2\sigma^{2} E_{\mathbb{Q}^{w_{u}},\mathbb{Q}^{w_{u}}}[|X-Y|\nu(X)] E_{\mathbb{Q}^{w_{u}},\mathbb{Q}^{w_{u}}}[|X-Y|\nu(X)]^{T}}{\sigma^{2} E_{\mathbb{Q}^{w_{u}},\mathbb{Q}^{w_{u}}}[|X-Y|]^{2}} \right), \end{split}$$

i.e.

$$s(\mathbb{Q}_{\theta}^{w_u}) = s(\mathbb{Q}_{(0,1)}^{w_u})$$

for $u \ge k$, showing that the swCRPS is locally tail-scale invariant. \Box

Proof of Proposition 3.8.

(i) According to Bolin and Wallin (2023), the scale function of rCRPS with kernel function

$$g_c(x, y) = \begin{cases} |x - y|, & |x - y| < c \\ 0, & \text{otherwise} \end{cases}$$
(C.5)

is

$$s(\mathbb{Q}_{\theta}) = \sigma^{-1} E_{\mathbb{Q},\mathbb{Q}} \left[H_{\mathbb{Q}}(X,Y) \mathbb{1} \left(|X-Y| < \frac{c}{\sigma} \right) \right],$$

where $H_{\mathbb{Q}}(X, Y)$ is a 2 × 2 positive semidefinite matrix independent of θ . Choose $u = \frac{c}{2\sigma}$. Then for |x| < u, |y| < u,

 $|x-y| \le |x| + |y| < \frac{c}{2\sigma},$ and

$$s\left(\mathbb{Q}_{\theta}^{w^{u}}\right) = \sigma^{-1} E_{\mathbb{Q}^{w^{u}},\mathbb{Q}^{w^{u}}}\left[H_{\mathbb{Q}^{w^{u}}}(X,Y)\right].$$

Therefore, rCRPS is not locally weight-scale invariant with respect to $w^u(x) = 1\{|x| < u\}$

(ii) Following the proof of Proposition 3.7, using the fact of Eq. (C.5) instead of Eq. (C.4) yields that the rSCRPS is locally weight-scale invariant for weight function $w(x) = 1\{|x| < c/2\}$. \Box

International Journal of Forecasting xxx (xxxx) xxx

Proof of Proposition 3.9. For $u \ge q$, the censored likelihood score, LS_q , and the logarithmic score, LS, have the same conditional expectation, since

$$S_u(\mathbb{P}, \mathbb{Q}) = E_{\mathbb{Q}}[LS_q(\mathbb{P}_u, Y)|Y > u]$$

= $E_{\mathbb{Q}}[1\{Y \le q\}\log(F(q))$
+ $1\{Y > q\}\log(f(Y))|Y > u]$
= $E_{\mathbb{Q}}[\log(f(Y))|Y > u]$
= $E_{\mathbb{Q}}[LS(\mathbb{P}_u, Y)|Y > u].$

The logarithmic score is locally scale invariant (Bolin & Wallin, 2023), and thus also locally tail-scale invariant according to Remark 3.2. Therefore, the conditional like-lihood is also locally tail-scale invariant. \Box

Proof of Proposition 3.10. First note that

$$CLog_w(\mathbb{P}, y) = w(y) \log\left(\int_{\mathbb{R}} p(t)w(t)dt\right)$$
$$+ (1 - w(y)) \log\left(1 - \int_{\mathbb{R}} p(t)w(t)dt\right)$$
$$+ w(y) \log(p_w(y))$$
$$= w(y) \log(p(y)) + (1 - w(y))$$
$$\times \log\left(1 - \int_{\mathbb{R}} p(t)w(t)dt\right)$$

Conditioned on $\tilde{w}(y) = 1\{w(y) = 1\}$, the score becomes the scale invariant logarithmic score, so the $CLog_w$ is locally weight-scale invariant with respect to \tilde{w} . \Box

Proof of Proposition 3.11.

(i) Consider the score

$$owS(\mathbb{P}, y; w) = w(y)S_0(\mathbb{P}^w, y)$$

Note that

$$owS(\mathbb{P}^w, \mathbb{Q}^w; w) = E[w(Y)S_0(\mathbb{P}^w, Y)|w(Y) > 0]$$

$$= E[S_0(\mathbb{P}^w, Y)|w(Y) > 0] = S_0(\mathbb{P}^w, \mathbb{Q}^w)$$

so if S₀ is locally weight-scale invariant with respect to w, then owS is also locally weight-scale invariant with respect to w.

(ii) Consider the score $S = S_1 + S_2$. The weight-scale function of *S* is the sum of the weight-scale functions of S_1 and S_2 . If S_1 and S_2 are locally weight-scale invariant with respect to *w*, their weight-scale functions fulfil the requirements for local scale invariance, and hence the weight-scale function of *S* does too. Therefore, *S* is locally weight-scale invariant. \Box

Appendix D. Estimated flood densities

The annual maximum water levels were modelled with a stationary GEV distribution and the estimated parameters shown in Table 3. Stations Michigan Huron (station 2) and Lake Superior (station 4) were used for simulations, since those stations suggested stationarity

H.K. Olafsdottir, H. Rootzén and D. Bolin

International Journal of Forecasting xxx (xxxx) xxx



Fig. D.10. Density of estimated GEV distributions for flood stations as described in Table 3, after shifting the x-axis by the estimated mode of each station.

and since the scale parameter of Lake Michigan Huron was almost twice the scale parameter of Lake Superior. All stations had negative shape parameter. In Fig. D.10 the estimated density function is compared at all five stations, where the x-axis has been shifted by each respective mode for readability.

References

- Allen, S., Ginsbourger, D., & Ziegel, J. (2022). Evaluating forecasts for high-impact events using transformed kernel scores.
- Baringhaus, L., & Franz, C. (2004). On a new multivariate two-sample test. Journal of Multivariate Analysis, 88(1), 190–206. http://dx.doi. org/10.1016/S0047-259X(03)00079-4.
- Bolin, D., & Simas, A. B. (2023). rSPDE: Rational approximations of fractional stochastic partial differential equations. https://CRAN.Rproject.org/package=rSPDE. R package version 2.3.3.
- Bolin, D., & Wallin, J. (2023). Local scale invariance and robustness of proper scoring rules. *Statistical Science*, 38(1), 140–159. http: //dx.doi.org/10.1214/22-STS864.
- Brehmer, J. R., & Strokorb, K. (2019). Why scoring functions cannot assess tail properties. *Electronic Journal of Statistics*, 13(2), 4015–4034. http://dx.doi.org/10.1214/19-EJS1622.
- Cameletti, M., Lindgren, F., Simpson, D., & Rue, H. (2013). Spatiotemporal modeling of particulate matter concentration through the SPDE approach. AStA. Advances in Statistical Analysis, 97, 109–131.
- Coles, S. (2001). An introduction to statistical modeling of extreme values. Springer.
- Cugerone, K., & De Michele, C. (2015). Johnson SB as general functional form for raindrop size distribution. Water Resources Research, 51(8), 6276–6289. http://dx.doi.org/10.1002/2014WR016484.
- Dawid, A. P. (2007). The geometry of proper scoring rules. Annals of the Institute of Statistical Mathematics, 59(1), 77–93. http://dx.doi. org/10.1007/s10463-006-0099-8.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. Journal of Business & Economic Statistics, 13(3), 253–263. http://dx. doi.org/10.1080/07350015.1995.10524599.
- Diks, C., Panchenko, V., & van Dijk, D. (2011). Likelihood-based scoring rules for comparing density forecasts in tails. *Journal of Econometrics*, 163(2), 215–230. http://dx.doi.org/10.1016/j.jeconom.2011.04. 001.
- Friederichs, P., & Thorarinsdottir, T. L. (2012). Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics*, 23(7), 579–594. http://dx.doi. org/10.1002/env.2176.

- Geneva: World Health Association (2021). WHO global air quality guidelines. Particulate matter (PM2.5 and PM10), ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide. URL https://www.who.int/publications/i/item/9789240034228.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378. http://dx.doi.org/10.1198/ 016214506000001437.
- Gneiting, T., & Ranjan, R. (2011). Comparing density forecasts using thresholdand quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, 29(3), 411–422. http://dx.doi.org/10.1198/jbes. 2010.08110.
- Haiden, T., Janousek, M., Vitart, F., Ferranti, L., & Prates, F. (2019). Evaluation of ECMWF forecasts, including the 2019 upgrade. ECMWF Technical Memoranda, 853(October), http://dx.doi.org/10. 21957/mlvapkke.
- Holzmann, H., & Klar, B. (2017). Focusing on regions of interest in forecast evaluation. The Annals of Applied Statistics, 11(4), 2404–2431. http://dx.doi.org/10.1214/17-AOAS1088.
- Ingebrigtsen, R., Lindgren, F., Steinsland, I., & Martino, S. (2015). Estimation of a non-stationary model for annual precipitation in southern Norway using replicates of the spatial field. *Spatial Statistics*, 14(Part C), 338–364. http://dx.doi.org/10.1016/j.spasta.2015.07.003.
- Krainski, E. (2023). The piemonte dataset example. https://eliaskrainski. github.io/INLAspacetime/articles/web/piemonte.html. (Accessed 09 November 2023).
- Lerch, S., Thorarinsdottir, T. L., Ravazzolo, F., & Gneiting, T. (2017). Forecaster's dilemma: Extreme events and forecast evaluation. *Statistical Science*, 32(1), 106–127. http://dx.doi.org/10.1214/16-STS588.
- Lindgren, F., Bakka, H., Bolin, D., Krainski, E., & Rue, H. (2023). A diffusion-based spatio-temporal extension of Gaussian Matérn fields.
- NOAA, Great Lakes Environmental Research Laboratory (2015). Water levels of the Great Lakes. URL https://www.glerl.noaa.gov/pubs/ brochures/lakelevels/archive/lakelevels_Feb_2015.pdf.
- NOAA National Centers for Environmental Information (2019). Climate at a glance: Regional time series. URL https://www.ncdc.noaa.gov/ cag/.
- NOAA's National Weather Service (2019). NOAA Atlas 14 Time Series Data. URL https://hdsc.nws.noaa.gov/hdsc/pfds/pfds_series.html.
- Olafsdottir, H. K., Rootzén, H., & Bolin, D. (2021). Extreme Rainfall Events in the Northeastern United States Become More Frequent with Rising Temperatures, but Their Intensity Distribution Remains Stable. *Journal of Climate*, *34*(22), 8863–8877. http://dx.doi.org/10. 1175/JCLI-D-20-0938.1.
- Opschoor, A., van Dijk, D., & van der Wel, M. (2017). Combining density forecasts using focused scoring rules. *Journal of Applied Econometrics*, 32(7), 1298–1313. http://dx.doi.org/10.1002/jae.2575.

H.K. Olafsdottir, H. Rootzén and D. Bolin

- Rue, H., Martino, S., & Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models using integrated nested Laplace approximations (with discussion). Journal of the Royal Statistical Society B, 71, 319–392.
- Taillardat, M., Fougères, A.-L., Naveau, P., & de Fondeville, R. (2022). Evaluating probabilistic forecasts of extremes using continuous ranked probability score distributions. *International Journal of Forecasting*, 39(3), 1448–1459. http://dx.doi.org/10.1016/j.ijforecast. 2022.07.003.
- Thorarinsdottir, T. L., & Schuhen, N. (2018). Chapter 6 verification: Assessment of calibration and accuracy. In S. Vannitsem, D. S. Wilks, & J. W. Messner (Eds.), *Statistical postprocessing of ensemble forecasts* (pp. 155–186). Elsevier, http://dx.doi.org/10.1016/B978-0-12-812372-0.00006-6.
- Tödter, J., & Ahrens, B. (2012). Generalization of the ignorance score: Continuous ranked version and its decomposition. *Monthly Weather Review*, 140(1948), 2005–2017. http://dx.doi.org/10.1175/MWR-D-11-00266.1.
- Vandeskog, S., Martino, S., Castro-Camilo, D., & Rue, H. (2022). Modelling sub-daily precipitation extremes with the blended generalised extreme value distribution. *Journal of Agricultural, Biological, and Environmental Statistics*, 27, 598–621. http://dx.doi.org/10.1007/ s13253-022-00500-7.
- Zamo, M., & 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. http://dx.doi.org/10.1007/s11004-017-9709-7.