

ORIGINAL ARTICLE

Isotonic distributional regression

Alexander Henzi¹ | Johanna F. Ziegel¹  | Tilmann Gneiting^{2,3} 

¹University of Bern, Bern, Switzerland

²Heidelberg Institute for Theoretical Studies, Heidelberg, Germany

³Karlsruhe Institute of Technology, Karlsruhe, Germany

Correspondence

Johanna F. Ziegel, University of Bern, Bern, Switzerland.

Email: johanna.ziegel@stat.unibe.ch

Abstract

Isotonic distributional regression (IDR) is a powerful non-parametric technique for the estimation of conditional distributions under order restrictions. In a nutshell, IDR learns conditional distributions that are calibrated, and simultaneously optimal relative to comprehensive classes of relevant loss functions, subject to isotonicity constraints in terms of a partial order on the covariate space. Non-parametric isotonic quantile regression and non-parametric isotonic binary regression emerge as special cases. For prediction, we propose an interpolation method that generalizes extant specifications under the pool adjacent violators algorithm. We recommend the use of IDR as a generic benchmark technique in probabilistic forecast problems, as it does not involve any parameter tuning nor implementation choices, except for the selection of a partial order on the covariate space. The method can be combined with subsample aggregation, with the benefits of smoother regression functions and gains in computational efficiency. In a simulation study, we compare methods for distributional regression in terms of the continuous ranked probability score (CRPS) and L_2 estimation error, which are closely linked. In a case study on raw and post-processed quantitative precipitation forecasts from a leading numerical weather prediction system, IDR is competitive with state of the art techniques.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2021 The Authors. Journal of the Royal Statistical Society: Series B (Statistical Methodology) published by John Wiley & Sons Ltd on behalf of Royal Statistical Society

KEYWORDS

conditional distribution estimation, monotonicity, probabilistic forecast, proper scoring rule, stochastic order, subagging, weather prediction

1 | INTRODUCTION

There is an emerging consensus in the transdisciplinary literature that regression analysis should be distributional, with Hothorn et al. (2014) arguing forcefully that

[t]he ultimate goal of regression analysis is to obtain information about the conditional distribution of a response given a set of explanatory variables.

Distributional regression marks a clear break from the classical view of regression, which has focused on estimating the conditional mean of the response variable in terms of one or more explanatory variable(s) or covariate(s). Later extensions have considered other functionals of the conditional distributions, such as quantiles or expectiles (Koenker, 2005; Newey & Powell, 1987; Schulze Waltrup et al., 2015). However, the reduction of a conditional distribution to a single-valued functional results in tremendous loss of information. Therefore, from the perspectives of both estimation and prediction, regression analysis ought to be distributional.

In the extant literature, both parametric and non-parametric approaches to distributional regression are available. Parametric approaches assume that the conditional distribution of the response is of a specific type (e.g. Gaussian) with an analytic relationship between the covariates and the distributional parameters. Key examples include statistically post-processed meteorological and hydrologic forecasts, as exemplified by Gneiting et al. (2005), Schefzik et al. (2013) and Vannitsem et al. (2018). In powerful semi-parametric variants, the conditional distributions remain parametric, but the influence of the covariates on the parameter values is modelled non-parametrically, for example by using generalized additive models (Klein et al., 2015; Rigby & Stasinopoulos, 2005; Umlauf & Kneib, 2018) or modern neural networks (Gasthaus et al., 2019; Rasp & Lerch, 2018). In related developments, semiparametric versions of quantile regression (Koenker, 2005) and transformation methods (Hothorn et al., 2014) can be leveraged for distributional regression.

Non-parametric approaches to distributional regression include kernel or nearest neighbour methods that depend on a suitable notion of distance on the covariate space. Then, the empirical distribution of the response for neighbouring covariates in the training set is used for distributional regression, with possible weighting in dependence on the distance to the covariate value of interest. Kernel smoothing methods and mixture approaches allow for absolutely continuous conditional distributions (Dunson et al., 2007; Hall et al., 1999; Li & Racine, 2008). Classification and regression trees partition the covariate space into leaves, and assign constant regression functions on each leaf (Breiman et al., 1984). Linear aggregation via bootstrap aggregation (bagging) or subsample aggregation (subagging) yields random forests (Breiman, 2001), which are increasingly being used to generate conditional predictive distributions, as proposed by Hothorn et al. (2004) and Meinshausen (2006).

Isotonicity is a natural constraint in estimation and prediction problems. Consider, for example, post-processing techniques in weather forecasting, where the covariates stem from the

output of numerical weather prediction (NWP) models, and the response variable is the respective future weather quantity. Intuitively, if the NWP model output indicates a larger precipitation accumulation, the associated regression functions ought to be larger as well. Isotonic relationships of this type hold in a plethora of applied settings. In fact, standard linear regression analysis rests on the assumption of isotonicity, in the form of monotonicity in the values of the covariate(s), save for changes in sign.

Concerning non-parametric regression for a conditional functional, such as the mean or a quantile, there is a sizable literature on estimation under the constraint of isotonicity. The classical work of Ayer et al. (1955), Bartholomew (1959a, b), Brunk (1955), van Eeden (1958), Miles (1959) is summarized in Barlow et al. (1972), de Leeuw et al. (2009), Robertson et al. (1988). Subsequent approaches include Bayesian and non-Bayesian smoothing techniques (e.g. Dette et al., 2006; Mammen, 1991; Neelon & Dunson, 2004; Shively et al., 2009), and reviews are available in Groeneboom & Jongbloed (2014) and Guntuboyina & Sen (2018).

In distributional regression, it may not be immediately clear what is meant by isotonicity, and the literature typically considers one ordinal covariate only (e.g. Davidov & Iliopoulos, 2012; El Barmi & Mukerjee, 2005; Hogg, 1965; Rojo & El Barmi, 2003), with a notable exception being the work of Mösching & Dümbgen (2020b), whose considerations allow for a real-valued covariate. In the general case of a partially ordered covariate space, which we consider here, it is unclear whether semi- or non-parametric techniques might be capable of handling monotonicity constraints, and suitable notions of isotonicity remain to be developed.

To this end, we assume that the response Y is real-valued, and equip the covariate space \mathcal{X} with a partial order \leq . Our aim is to estimate the conditional distribution of Y given the covariate X , for short $\mathcal{L}(Y|X)$, on training data, in a way that respects the partial order, and we desire to use this estimate for prediction. Formally, a distributional regression technique generates a mapping from $x \in \mathcal{X}$ to a probability measure F_x , which serves to model the conditional distribution $\mathcal{L}(Y|X = x)$. This mapping is isotonic if $x \leq x'$ implies $F_x \leq_{\text{st}} F_{x'}$, where \leq_{st} denotes the usual stochastic order, that is, $G \leq_{\text{st}} H$ if $G(y) \geq H(y)$ for $y \in \mathbb{R}$, where we use the same symbols for the probability measures G, H and their associated conditional cumulative distribution functions (CDFs). Equivalently, $G \leq_{\text{st}} H$ holds if $G^{-1}(\alpha) \leq H^{-1}(\alpha)$ for $\alpha \in (0, 1)$, where $G^{-1}(\alpha) = \inf\{y \in \mathbb{R}: G(y) \geq \alpha\}$ is the standard quantile function (Shaked & Shanthikumar, 2007).

Useful comparisons of predictive distributions are in terms of proper scoring rules, of which the most prominent and most relevant instance is the continuous ranked probability score (CRPS; Gneiting & Raftery, 2007; Matheson & Winkler, 1976). We show that there is a unique isotonic distributional regression that is optimal with respect to the CPRS (Theorem 1), and refer to it as the *isotonic distributional regression* (IDR). As it turns out, IDR is a universal solution, in that the estimate is optimal with respect to a broad class of proper scoring rules (Theorem 2). Classical special cases such as non-parametric isotonic quantile regression and probabilistic classifiers for threshold-defined binary events are nested by IDR. Simultaneously, IDR avoids pitfalls commonly associated with non-parametric distributional regression, such as suboptimal partitions of the covariate space and level crossing (Athey et al., 2019, p. 1167).

For illustration, consider the joint distribution of (X, Y) , where X is uniform on $(0, 10)$ and

$$Y|X \sim \text{Gamma}(\text{shape} = \sqrt{X}, \text{scale} = \min\{\max\{X, 1\}, 6\}), \quad (1)$$

so that $\mathcal{L}(Y|X = x) \leq_{\text{st}} \mathcal{L}(Y|X = x')$ if $x \leq x'$. Figure 1 shows IDR conditional CDFs and quantiles as estimated on a training set of size $n = 600$. IDR is capable of estimating both the strongly right-skewed conditional distributions for lower values of X and the more symmetric distributions as X




FIGURE 1 Simulation example for a sample of size $n = 600$ from the distribution in (1): (a) True conditional CDFs (smooth) and IDR estimates (step functions) for selected values of the covariate. (b) IDR estimated conditional distributions. The shaded bands correspond to probability mass 0.10 each, with the darkest shade marking the central interval. Vertical strips indicate the cross-sections corresponding to the values of the covariate in panel (a) [Colour figure can be viewed at wileyonlinelibrary.com]

increases. The CDFs are piecewise constant, and they never cross each other. The computational cost of IDR is of order at least $\mathcal{O}(n \log n)$ and may become prohibitive as n grows. However, IDR can usefully be combined with subsample aggregation (subagging), much in the spirit of random forests (Breiman, 2001), with the benefits of reduced computational cost under large training samples, smoother regression functions, and (frequently) improved predictive performance.

The remainder of the paper is organized as follows. The methodological core of the paper is in Section 2, where we prove existence, uniqueness and universality of the IDR solution, discuss computational issues and asymptotic consistency, and propose strategies for prediction. In Section 3, we turn to the critical issue of the choice of a partial order on the covariate space. Section 4 reports on a comparative simulation study that addresses both prediction and estimation, and Section 5 is devoted to a case study on probabilistic quantitative precipitation forecasts, with covariates provided by the European Centre for Medium-Range Weather Forecasts (ECMWF) ensemble system. Precipitation accumulations feature unfavourable properties that challenge parametric approaches to distributional regression: The conditional distributions have a point mass at zero, and they are continuous and right skewed on the positive half-axis. In a

comparison to state-of-the-art methods that have been developed specifically for the purpose, namely Bayesian model averaging (BMA; Sloughter et al., 2007), ensemble model output statistics (EMOS; Scheuerer, 2014) and heteroscedastic censored logistic regression (HCLR; Messner et al., 2014), the (out-of-sample) predictive performance of IDR is competitive, despite the method being generic, and being fully automatic once a partial order on the covariate space has been chosen.

We close the paper with a discussion in Section 6, where we argue that IDR provides a very widely applicable, competitive benchmark in probabilistic forecasting problems. The use of benchmark techniques has been called for across application domains (e.g. Basel Committee on Banking Supervision, 2016; Pappenberger et al., 2015; Rossi, 2013; Vogel et al., 2018), and suitable methods should be competitive in terms of predictive performance, while avoiding implementation decisions that may vary from user to user. IDR is well suited to this purpose, as it is entirely generic, does not involve any implementation decisions, other than the choice of the partial order, applies to all types of real-valued outcomes with discrete, continuous or mixed discrete-continuous distributions, and accommodates general types of covariate spaces.

2 | ISOTONIC DISTRIBUTIONAL REGRESSION

We proceed to introduce the IDR technique. To this end, we first review basic facts on proper scoring rules and notions of calibration. Then we define the IDR solution, prove existence, uniqueness and universality, and discuss its computation and asymptotic consistency. Thereafter, we turn from estimation to prediction and describe how IDR can be used in out-of-sample forecasting. Throughout, we identify a Borel probability measure on the real line \mathbb{R} with its cumulative distribution function (CDF), and we denote the extended real line by $\overline{\mathbb{R}} = [-\infty, \infty]$.

2.1 | Preliminaries

Following Gneiting & Raftery (2007), we argue that distributional regression techniques should be compared and evaluated using proper scoring rules. A *proper scoring rule* is a function $S: \mathcal{P} \times \mathbb{R} \rightarrow \overline{\mathbb{R}}$, where \mathcal{P} is a suitable class of probability measures on \mathbb{R} , such that $S(F, \cdot)$ is measurable for any $F \in \mathcal{P}$, the integral $\int S(G, y) dF(y)$ exists, and

$$\int S(F, y) dF(y) \leq \int S(G, y) dF(y)$$

for all $F, G \in \mathcal{P}$. A key example is the *continuous ranked probability score* (CRPS), which is defined for all Borel probability measures, and given as

$$\text{CRPS}(F, y) = \int_{\mathbb{R}} (F(z) - \mathbb{1}_{\{y \leq z\}})^2 dz.$$

Introduced by Matheson & Winkler (1976), the CRPS has become popular across application areas and methodological communities, both for the purposes of evaluating predictive performance and as a loss function in estimation; see, for example, Gasthaus et al. (2019), Gneiting et al. (2005), Hersbach (2000), Hothorn et al. (2014), Pappenberger et al. (2015) and Rasp & Lerch (2018). The

CRPS is reported in the same unit as the response variable, and it reduces to the absolute error, $|x - y|$, if F is the point or Dirac measure in $x \in \mathbb{R}$.

Results in Ben Bouallègue et al. (2018), Ehm et al. (2016) and Laio & Tamea (2007) imply that the CRPS can be represented equivalently as

$$\text{CRPS}(F, y) = 2 \int_{(0,1)} \text{QS}_\alpha(F, y) d\alpha \quad (2)$$

$$= 2 \int_{(0,1)} \int_{\mathbb{R}} S_{\alpha, \theta}^Q(F, y) d\theta d\alpha \quad (3)$$

$$= \int_{\mathbb{R}} \int_{(0,1)} S_{z, c}^P(F, y) dc dz, \quad (4)$$

where the mixture representation (2) is in terms of the asymmetric piecewise linear or pinball loss,

$$\text{QS}_\alpha(F, y) = \begin{cases} (1 - \alpha)(F^{-1}(\alpha) - y), & y \leq F^{-1}(\alpha), \\ \alpha(y - F^{-1}(\alpha)), & y \geq F^{-1}(\alpha), \end{cases} \quad (5)$$

which is customarily thought of as a quantile loss function, but can be identified with a proper scoring rule (Gneiting, 2011, Theorem 3). The representations (3) and (4) express the CRPS in terms of the *elementary* or *extremal scoring functions* for the α -quantile functional, namely,

$$S_{\alpha, \theta}^Q(F, y) = \begin{cases} 1 - \alpha, & y \leq \theta < F^{-1}(\alpha), \\ \alpha, & F^{-1}(\alpha) \leq \theta < y, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where $\theta \in \mathbb{R}$; and for probability assessments of the binary outcome $\mathbb{1}\{y \leq z\}$ at the threshold value $z \in \mathbb{R}$, namely

$$S_{z, c}^P(F, y) = \begin{cases} 1 - c, & F(z) < c, \quad y \leq z, \\ c, & F(z) \geq c, \quad y > z, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where $c \in (0, 1)$. For background information on elementary or extremal scoring functions and related concepts see Ehm et al. (2016).

Predictive distributions ought to be calibrated (Dawid, 1984; Diebold et al., 1998; Gneiting et al., 2007), in the broad sense that they should be statistically compatible with the responses, and various notions of calibration have been proposed and studied. In the spirit of Gneiting & Ranjan (2013), we consider the joint distribution \mathbb{P} of the response Y and the distributional regression F_X . The most widely used criterion is *probabilistic calibration*, which requires that the *probability integral transform* (PIT), namely, the random variable

$$Z = F_X(Y-) + V(F_X(Y) - F_X(Y-)), \quad (8)$$

be standard uniform, where $F_X(Y-) = \lim_{y \uparrow Y} F_X(y)$ and V is a standard uniform variable that is independent of F_X and Y . If F_X is continuous the PIT is simply $Z = F_X(Y)$. Here we introduce the novel notion of *threshold calibration*, requiring that

$$\mathbb{P}(Y \leq y | F_X(y)) = F_X(y) \quad (9)$$

almost surely for $y \in \mathbb{R}$, which implies *marginal calibration*, defined as $\mathbb{P}(Y \leq y) = \mathbb{E}(F_X(y))$ for $y \in \mathbb{R}$. If $F_X = \mathcal{L}(Y|X)$ then it is calibrated in any of the above senses (Gneiting & Ranjan, 2013, Theorem 2.8).

2.2 | Existence, uniqueness and universality

A partial order relation \leq on a set \mathcal{X} has the same properties as a total order, namely reflexivity, antisymmetry and transitivity, except that the elements need not be comparable, that is, there might be elements $x \in \mathcal{X}$ and $x' \in \mathcal{X}$ such that neither $x \leq x'$ nor $x' \leq x$ holds. A key example is the componentwise order on \mathbb{R}^d .

For a positive integer n and a partially ordered set \mathcal{X} , we define the classes

$$\begin{aligned} \mathcal{X}_{\uparrow}^n &= \{\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n: x_1 \leq \dots \leq x_n\}, \\ \mathcal{X}_{\downarrow}^n &= \{\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n: x_1 \geq \dots \geq x_n\} \end{aligned}$$

of the increasingly and decreasingly (totally) ordered tuples in \mathcal{X} , respectively. Similarly, given a further partially ordered set \mathcal{Q} and a vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$, the classes

$$\begin{aligned} \mathcal{Q}_{\uparrow, \mathbf{x}}^n &= \{\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{Q}^n: q_i \leq q_j \text{ if } x_i \leq x_j\}, \\ \mathcal{Q}_{\downarrow, \mathbf{x}}^n &= \{\mathbf{q} = (q_1, \dots, q_n) \in \mathcal{Q}^n: q_i \geq q_j \text{ if } x_i \leq x_j\} \end{aligned}$$

comprise the increasingly and decreasingly (partially) ordered tuples in \mathcal{Q} , with the order induced by the tuple \mathbf{x} and the partial order \leq on \mathcal{X} .

Let $I \subseteq \mathbb{R}$ be an interval, and let S be a proper scoring rule with respect to a class \mathcal{P} of probability distributions on I that contains all distributions with finite support. Given training data in the form of a covariate vector $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$ and response vector $\mathbf{y} = (y_1, \dots, y_n) \in I^n$, we may interpret any mapping from $\mathbf{x} \in \mathcal{X}^n$ to \mathcal{P}^n as a distributional regression function. Throughout, we equip \mathcal{P} with the usual stochastic order.

Definition 1 (S-based regression). An element $\hat{\mathbf{F}} = (\hat{F}_1, \dots, \hat{F}_n) \in \mathcal{P}^n$ is an *S-based isotonic regression* of $\mathbf{y} \in I^n$ on $\mathbf{x} \in \mathcal{X}^n$, if it is a minimizer of the empirical loss

$$\ell_S(\mathbf{F}) = \frac{1}{n} \sum_{i=1}^n S(F_i, y_i)$$

over all $\mathbf{F} = (F_1, \dots, F_n)$ in $\mathcal{P}_{\uparrow, \mathbf{x}}^n$.

In plain words, an S-based isotonic regression achieves the best fit in terms of the scoring rule S , subject to the conditional CDFs $\hat{F}_1, \dots, \hat{F}_n$ satisfying partial order constraints induced by the covariate values x_1, \dots, x_n . The definition and the subsequent results can be extended to losses of

the form $\ell_S(\mathbf{F}) = \sum_{i=1}^n w_i S(F_i, y_i)$ with rational, strictly positive weights w_1, \dots, w_n . The adaptations are straightforward and left to the reader.

Furthermore, the definition of an S-based isotonic regression as a minimizer of ℓ_S continues to apply when \mathcal{X} is equipped with a pre- or quasiorder \leq instead of a partial order. Preorders are not necessarily antisymmetric, and so there might be elements x, x' such that $x \leq x'$ and $x' \leq x$ but $x' \neq x$. In this setting, we define x and x' to be equivalent if $x \leq x'$ and $x' \leq x$, and set $[x] \leq_p [x']$ if representatives u, u' of the equivalence classes $[x], [x']$ satisfy $u \leq u'$. Then the binary relation \leq_p defines a partial order on the set of equivalence classes, and the S-based isotonic regression with the new covariates and the partial order \leq_p coincides with the original S-based isotonic regression.

In Supplementary Section S1 we prove the following result.

Theorem 1 (existence and uniqueness). *There exists a unique CRPS-based isotonic regression $\hat{\mathbf{F}} \in \mathcal{P}^n$ of \mathbf{y} on \mathbf{x} .*

We refer to this unique $\hat{\mathbf{F}}$ as the IDR of \mathbf{y} on \mathbf{x} . In the particular case of a total order on the covariate space, and assuming that $x_1 < \dots < x_n$, for each $z \in I$ the solution $\hat{\mathbf{F}}(z) = (\hat{F}_1(z), \dots, \hat{F}_n(z))$ is given by

$$\hat{F}_i(z) = \min_{k=1, \dots, i} \max_{j=k, \dots, n} \frac{1}{j-k+1} \sum_{l=k}^j \mathbb{1}\{y_l \leq z\} \quad (10)$$

for $i = 1, \dots, n$; see Equations (1.9)–(1.13) of Barlow et al. (1972). A similar max–min formula applies under partial orders (Jordan et al., 2021; Robertson & Wright, 1980), and it follows that \hat{F}_i is piecewise constant with any points of discontinuity at y_1, \dots, y_n .

At first sight, the specific choice of the CRPS as a loss function may seem arbitrary. However, the subsequent result, which we prove in Supplementary Section S1, reveals that IDR is simultaneously optimal with respect to broad classes of proper scoring rules that include all relevant choices in the extant literature. The popular logarithmic score allows for the comparison of absolutely continuous distributions with respect to a fixed dominating measure only and thus is not applicable here. Statements concerning calibration are with respect to the empirical distribution of the training data $(x_1, y_1), \dots, (x_n, y_n)$.

Theorem 2 (universality). *The IDR solution $\hat{\mathbf{F}}$ of \mathbf{y} on \mathbf{x} is threshold calibrated and has the following properties.*

- (a) *The IDR solution $\hat{\mathbf{F}}$ is an S-based isotonic regression of y on x under any scoring rule of the form*

$$S(F, y) = \int_{(0, 1) \times \mathbb{R}} S_{\alpha, \theta}^Q(F, y) dH(\alpha, \theta) \quad (11)$$

or

$$S(F, y) = \int_{\mathbb{R} \times (0, 1)} S_{z, c}^P(F, y) dM(z, c), \quad (12)$$

where $S_{\alpha,\theta}^Q$ is the elementary quantile scoring function (6), $S_{z,c}^P$ is the elementary probability scoring rule (7), and H and M are locally finite Borel measures on $(0, 1) \times \mathbb{R}$ and $\mathbb{R} \times (0, 1)$, respectively.

(b) For every $\alpha \in (0, 1)$ it holds that $\hat{\mathbf{F}}^{-1}(\alpha) = (\hat{F}_1^{-1}(\alpha), \dots, \hat{F}_n^{-1}(\alpha))$ is a minimizer of

$$\frac{1}{n} \sum_{i=1}^n s_{\alpha}(\theta_i, y_i) \quad (13)$$

over all $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n) \in I_{\uparrow, \mathbf{x}}^n$, under any function $s_{\alpha}: I \times I \rightarrow \overline{\mathbb{R}}$ which is left-continuous in both arguments and such that $S(F, y) = s_{\alpha}(F^{-1}(\alpha), y)$ is a proper scoring rule on \mathcal{P} .

(c) For every threshold value $z \in I$, it is true that $\hat{\mathbf{F}}(z) = (\hat{F}_1(z), \dots, \hat{F}_n(z))$ is a minimizer of

$$\frac{1}{n} \sum_{i=1}^n s(\eta_i, \mathbb{1}_{\{y_i \leq z\}}) \quad (14)$$

over all ordered tuples $\boldsymbol{\eta} = (\eta_1, \dots, \eta_n) \in [0, 1]_{\downarrow, \mathbf{x}}^n$, under any function $s: [0, 1] \times \{0, 1\} \rightarrow \overline{\mathbb{R}}$ that is a proper scoring rule for binary events, which is left-continuous in its first argument, satisfies $s(0, y) = \lim_{p \rightarrow 0} s(p, y)$, and is real-valued, except possibly $s(0, 1) = -\infty$ or $s(1, 0) = -\infty$.

The quantile weighted and threshold weighted versions of the CRPS studied by Gneiting & Ranjan (2011) arise from (11) and (12) with $H = G_0 \otimes \lambda$ and $M = \lambda \otimes G_1$, where λ denotes the Lebesgue measure, and G_0 and G_1 are σ -finite Borel measures on $(0, 1)$ and \mathbb{R} , respectively. If G_0 and G_1 are Lebesgue measures, we recover the mixture representations (3) and (4) of the CRPS. By results of Ehm et al. (2016), if H is concentrated on $\{\alpha\} \times \mathbb{R}$ and M is concentrated on $\{z\} \times (0, 1)$, these representations cover essentially all proper scoring rules that depend on the predictive distribution F via $F^{-1}(\alpha)$ or $F(z)$ only, yielding universal optimality in statements in parts (b) and (c) of Theorem 2.

In particular, as a special case of (13), the IDR solution is a minimizer of the quantile loss under the asymmetric piecewise linear or pinball function (5) that lies at the heart of quantile regression (Koenker, 2005). Consequently, as the mixture representation (2) of the CRPS may suggest, IDR nests classical non-parametric isotonic quantile regression as introduced and studied by Casady & Cryer (1976) and Robertson & Wright (1975). In other words, part (b) of Theorem 2 demonstrates that, if we (hypothetically) perform non-parametric isotonic quantile regression at every level $\alpha \in (0, 1)$ and piece the conditional quantile functions together, we recover the IDR solution. However, the IDR solution is readily computable (Section 2.3), without invoking approximations or truncations, unlike brute force approaches to simultaneous quantile regressions. Loss functions of the form (13) also include the interval score (Gneiting & Raftery, 2007, equation (43); Winkler, 1972), which constitutes the most used proper performance measure for interval forecasts.

In the special case of a binary response variable, we see from (c) and (14) that the IDR solution is an S-based isotonic regression under just any applicable proper scoring rule S . Furthermore, threshold calibration is the strongest possible notion of calibration in this setting (Gneiting & Ranjan, 2013, Theorem 2.11), so the IDR solution is universal in every regard. In the further special case of a total order on the covariate space, the IDR and pool adjacent violators (PAV) algorithm solutions coincide, and the statement in (c) is essentially equivalent to Theorem 1.12 of Barlow et al. (1972). In particular, the IDR or PAV solution yields both the non-parametric maximum likelihood estimate and the non-parametric least squares estimate under the constraint of isotonicity. The latter suggests a computational implementation via quadratic programming, to which we tend now.

2.3 | Computational aspects

The key observation towards a computational implementation is the aforementioned special case of (14), according to which the IDR solution $\hat{\mathbf{F}} \in \mathcal{P}^n$ of $\mathbf{y} \in \mathbb{R}^n$ on $\mathbf{x} \in \mathcal{X}^n$ satisfies

$$\hat{\mathbf{F}}(z) = \arg \min_{\eta \in [0,1]_{\downarrow, \mathbf{x}}^n} \sum_{i=1}^n (\eta_i - \mathbb{1}\{y_i \leq z\})^2 \quad (15)$$

at every threshold value $z \in \mathbb{R}$. In this light, the computation of the IDR CDF at any fixed threshold reduces to a quadratic programming problem. The above target function is constant in between the unique values of y_1, \dots, y_n , say $\tilde{y}_1 < \dots < \tilde{y}_m$, and so it suffices to estimate the CDFs at these points only. In contrast, exact implementations based on quantiles would need to consider all levels of the form i/j with integers $1 \leq i < j \leq n$, which is computationally prohibitive. In the threshold-based approach, the overall cost depends on the quadratic programming solver applied, and the computation becomes much faster if recursive relations between consecutive conditional CDFs $\hat{\mathbf{F}}(\tilde{y}_k)$ and $\hat{\mathbf{F}}(\tilde{y}_{k-1})$ are taken advantage of. In the case of a total order, Henzi et al. (2020) describe a recursive adaptation of the PAV algorithm to IDR that considerably reduces the computation time as compared to a naive implementation which does not take into account recursive relations. Under general partial orders, active set methods for solutions to the quadratic programming problem (15) have been discussed by de Leeuw et al. (2009). In our implementation, we use the powerful quadratic programming solver OSQP (Stellato et al., 2020) as supplied by the package `osqp` in the statistical programming environment R (R Core Team, 2020; Stellato et al., 2019), which can be warm-started efficiently by taking $\hat{\mathbf{F}}(\tilde{y}_{k-1})$ as a starting point for the computation of $\hat{\mathbf{F}}(\tilde{y}_k)$.

Clearly, a challenge in the computational implementation of IDR with general partial orders is that the number of variables in the quadratic programming problem (15) grows at a rate of $\mathcal{O}(n)$. As a remedy, we propose subsample aggregation, much in the spirit of random forests that rely on bootstrap aggregated (bagged) classification and regression trees (Breiman, 1996, 2001). It was observed early on that random forests generate conditional predictive distributions (Hothorn et al., 2004; Meinshausen, 2006), and recent applications include the statistical post-processing of ensemble weather forecasts (Schlosser et al., 2019; Taillardat et al., 2016; Taillardat et al., 2019). Bühlmann & Yu (2002) and Buja & Stützel (2006) argue forcefully that subsample aggregation (subagging) tends to be equally effective as bagging, but at considerably lower computational cost.

In view of the superlinear computational costs of IDR, smart uses of subsample aggregation yield major efficiency gains, taking into account that the estimation on different subsamples can be performed in parallel. Isotonicity is preserved under linear aggregation, and the aggregated conditional CDFs can be inverted to generate isotonic conditional quantile functions, with the further benefit of smoother estimates in continuous settings. A detailed investigation of optimal subsample aggregation for IDR is a topic for future research. For illustration, Figure 2 returns to the simulation example in Figure 1, but now with a much larger training sample of size $n = 10,000$ from the distribution in (1). Linear aggregation based on 100 subsamples (drawn without replacement) of size $n = 1000$ each is superior to the brute force approach on the full training set in terms of estimation accuracy. The computation on the full dataset for this simulation example takes 11.7 s for the naive implementation, but only 1.1 s for the sequential algorithm of Henzi et al. (2020). Subagging gives computation times of 9.9 and 2.5 s, respectively, or 1.8 and 0.5 s when parallelized over eight cores.¹

¹With Intel(R) Xeon(R) E5-2630 v4 2.20GHz CPUs, in R (R Core Team, 2020), using the `doParallel` package for parallelization. Times reported are averages over 100 replicates.




FIGURE 2 Simulation example for a sample of size $n = 10,000$ from the distribution in (1). The true conditional CDFs (smooth dashed graphs) are compared to IDR estimates (step functions) based on (a) the full training sample of size $n = 10,000$ and (b) linear aggregation of IDR estimates on 100 subsamples of size 1000 each [Colour figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com)]

2.4 | Consistency

We proceed to prove uniform consistency of the IDR estimator. While strong consistency of non-parametric isotonic quantile regression for single quantiles was proved decades ago (Casady & Cryer, 1976; Robertson & Wright, 1975), uniform consistency and rates of convergence for the IDR estimator have been established only recently, and exclusively in the case of a total order, see El Barmi & Mukerjee (2005, Theorem 1) and Mösching & Dümbgen (2020b, Theorem 3.3).

For $x \in \mathcal{X}$ and $y \in \mathbb{R}$, let $\hat{F}_x(y)$ denote the IDR estimate based on fixed or random pairs $(X_1, Y_1), \dots, (X_n, Y_n)$. As introduced thus far, the IDR solution $\hat{F} = (\hat{F}_1, \dots, \hat{F}_n)$ is defined at the covariate values $X_1, \dots, X_n \in \mathcal{X}$ only. For general $x \in \mathcal{X}$, we merely assume that $\hat{F}_x(y)$ is some value in between the bounds given by

$$\max_{i \in s(x)} \hat{F}_i(y) \leq \hat{F}_x(y) \leq \min_{i \in p(x)} \hat{F}_i(y). \quad (16)$$

Here, we define the sets of the indices of *direct predecessors* and *direct successors* of $x \in \mathcal{X}$ among the covariate values as

$$p(x) = \{i \in \{1, \dots, n\} : X_i \leq X_j \leq x \Rightarrow X_j = X_i, j = 1, \dots, n\} \quad (17)$$

and

$$s(x) = \{i \in \{1, \dots, n\} : x \leq X_j \leq X_i \Rightarrow X_j = X_i, j = 1, \dots, n\}, \quad (18)$$

respectively.

In Supplementary Section S2, we establish the following consistency theorem, which covers key examples of partial orders and is based on strictly weaker assumptions than the results of Mösching & Dümbgen (2020b). However, in contrast to their work, we do not provide rates of convergence. The choice $\mathcal{X} = [0, 1]^d$ for the covariate space merely serves to simplify the presentation: As IDR is invariant under strictly isotonic transformations, any covariate vector $X = (X_1, \dots, X_d) \in \mathbb{R}^d$ can be transformed to have support in $[0, 1]^d$, and the componentwise partial order can be replaced by any weaker preorder. A key assumption uses the concept of an *antichain* in a partially ordered set (S, \leq) , which is a subset $A \subseteq S$ that does not admit comparisons, in the sense that $u \leq v$ for $u, v \in A$ implies $u = v$. As we discuss subsequently, results of Brightwell (1992) imply that the respective distributional condition is mild.

Theorem 3 (uniform consistency). *Let $\mathcal{X} = [0, 1]^d$ be endowed with the componentwise partial order and the norm $\|u\| = \max_{i=1, \dots, d} |u_i|$. Let further $(X_{ni}, Y_{ni}) \in [0, 1]^d \times \mathbb{R}$, $n \in \mathbb{N}$, $i = 1, \dots, n$, be a triangular array such that $(X_{n1}, Y_{n1}), \dots, (X_{nn}, Y_{nn})$ are independent and identically distributed random vectors for each $n \in \mathbb{N}$, and let $S_n = \{X_{n1}, \dots, X_{nn}\}$. Assume that*

(a) *for all non-degenerate rectangles $J \subseteq \mathcal{X}$, there exists a constant $c_J > 0$ such that*

$$\#(S_n \cap J) \geq nc_J$$

with asymptotic probability one, that is, if A_n denotes the event that $\#(S_n \cap J) \geq nc_J$, then $\mathbb{P}(A_n) \rightarrow 1$ as $n \rightarrow \infty$;

(b) *for some $\gamma \in (0, 1)$,*

$$\max\{\#A : A \subset S_n \text{ is antichain}\} \leq n^\gamma$$

with asymptotic probability one.

Assume further that the true conditional CDFs $F_x(y) = \mathbb{P}(Y_{ni} \leq y | X_{ni} = x)$ satisfy

(c) *$F_x(y)$ is decreasing in x for all $y \in \mathbb{R}$;*

(d) *for every $\eta > 0$, there exists $r > 0$ such that*

$$\sup\{|F_x(y) - F_{x'}(y)| : x, x' \in [0, 1]^d, \|x - x'\| \leq r, y \in \mathbb{R}\} < \eta.$$

Then for every $\varepsilon > 0$ and $\delta > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{x \in [\delta, 1-\delta]^d, y \in \mathbb{R}} |\hat{F}_x(y) - F_x(y)| \geq \varepsilon \right) = 0. \quad (19)$$

Assumption (a) requires that the covariates are sufficiently dense in \mathcal{X} , as is satisfied under strictly positive Lebesgue densities on \mathcal{X} . In order to derive rates of convergence, the size of the rectangles J in (a) would need to decrease with n , as in condition (A.2) of Mösching & Dümbgen (2020b); we leave this type of extension as a direction for future work. Assumption (c) is the basic model assumption of IDR, while assumption (d) requires uniform continuity of the conditional distributions, which is weaker than Hölder continuity in condition (A.1) of Mösching & Dümbgen (2020b).

Assumption (b), which is always satisfied in the case of a total order, calls for a more detailed discussion. In words, the maximal number of mutually incomparable elements needs to grow at a rate slower than n^γ . Evidently, the easier elements can be ordered, the smaller the maximal antichain. Consequently, Theorem 3 continues to hold under the empirical stochastic order and the empirical increasing convex order on the covariates introduced in Section 3.3, and indeed under any preorder that is weaker than the componentwise order. The key to understanding the distributional implications of (b) is Corollary 2 in Brightwell (1992), which states that for a sequence of independent random vectors from a uniform population on $[0, 1]^d$ the size of a maximal antichain grows at a rate of $n^{1-1/d}$; see also the remark following the proof of Theorem 3 in Supplementary Section S2.

As comparability under the componentwise order is preserved under monotonic transformations, any covariate vector $X \in \mathbb{R}^d$ that can be obtained as a monotone transformation of a uniform random vector of arbitrary dimension guarantees (b). This includes, for example all Gaussian random vectors with non-negative correlation coefficients. In this light, assumption (b) is rather weak, and well in line with the intuition that for multivariate isotonic (distributional) regression to work well, there ought to be at least minor positive dependence between the covariates. In the context of our case study in Section 5, high positive correlations between the covariates are the rule, as exemplified by Table 3 in Raftery et al. (2005).

2.5 | Prediction

As noted, the IDR solution $\hat{\mathbf{F}} = (\hat{F}_1, \dots, \hat{F}_n) \in \mathcal{P}_{\uparrow, \mathbf{x}}^n$ is defined at the covariate values $x_1, \dots, x_n \in \mathcal{X}$ only. Generally, if a (not necessarily optimal) distributional regression $\mathbf{F} = (F_1, \dots, F_n) \in \mathcal{P}_{\uparrow, \mathbf{x}}^n$ is available, a key task in practice is to make a prediction at a new covariate value $x \in \mathcal{X}$ where $x \notin \{x_1, \dots, x_n\}$. We denote the respective predictive CDF by F .

In the specific case $\mathcal{X} = \mathbb{R}$ of a single real-valued covariate there is a simple way of doing this, as frequently implemented in concert with the PAV algorithm. For simplicity, we suppose that $x_1 < \dots < x_n$. If $x < x_1$ we may let $F = F_1$; if $x \in (x_i, x_{i+1})$ for some $i \in \{1, \dots, n-1\}$ we may interpolate linearly, so that

$$F(z) = \frac{x - x_i}{x_{i+1} - x_i} F_i(z) + \frac{x_{i+1} - x}{x_{i+1} - x_i} F_{i+1}(z)$$

for $z \in \mathbb{R}$, and if $x > x_n$ we may set $F = F_n$. However, approaches that are based on interpolation do not extend to a generic covariate space, which may or may not be equipped with a metric.

In contrast, the method we describe now, which generalizes a proposal by Wilbur et al. (2005), solely uses information supplied by the partial order \leq on the covariate space \mathcal{X} . For a general covariate value $x \in \mathcal{X}$, the sets of the indices of direct predecessors and direct successors among the covariate values x_1, \dots, x_n in the training data are defined as at (17) and (18), respectively, with X_1, \dots, X_n replaced by x_1, \dots, x_n . If the covariate space \mathcal{X} is totally ordered, these sets contain at most one element. If the order is partial but not total, $p(x)$ and $s(x)$ may, and frequently do, contain more than one element. Assuming that $p(x)$ and $s(x)$ are non-empty, any predictive CDF F that is consistent with \mathbf{F} must satisfy

$$\max_{i \in s(x)} F_i(z) \leq F(z) \leq \min_{i \in p(x)} F_i(z) \quad (20)$$

at all threshold values $z \in \mathbb{R}$. We now let F be the pointwise arithmetic average of these bounds, that is,

$$F(z) = \frac{1}{2} \left(\max_{i \in s(x)} F_i(z) + \min_{i \in p(x)} F_i(z) \right) \quad (21)$$

for $z \in \mathbb{R}$. If $s(x)$ is empty while $p(x)$ is non-empty, or vice versa, a natural choice, which we employ hereinafter, is to let F equal the available bound given by the non-empty set. If x is not comparable to any of x_1, \dots, x_n the training data lack information about the conditional distribution at x , and a natural approach, which we adopt and implement, is to set F equal to the empirical distribution of the response values y_1, \dots, y_n .

The difference between the bounds (if any) in (20) might be a useful measure of estimation uncertainty and could be explored as a promising avenue towards the quantification of ambiguity and generation of second-order probabilities (Ellsberg, 1961; Seo, 2009). In the context of ensemble weather forecasts, the assessment of ambiguity has been pioneered by Allen & Eckel (2012). Interesting links arise when the envelope in (20) is interpreted in the spirit of randomized predictive systems and conformal estimates as studied by Vovk et al. (2019); compare, for example, their Figure 5 with our Figure 4b below.

3 | PARTIAL ORDERS

The choice of a sufficiently informative partial order on the covariate space is critical to any successful application of IDR. In the extreme case of distinct, totally ordered covariate values $x_1, \dots, x_n \in \mathcal{X}$ and a perfect monotonic relationship to the response values y_1, \dots, y_n , the IDR distribution associated with x_i is simply the point measure in y_i , for $i = 1, \dots, n$. The same happens in the other extreme, when there are no order relations at all. Hence, the partial order serves to regularize the IDR solution.

Thus far, we have simply assumed that the covariate space \mathcal{X} is equipped with a partial order \leq , without specifying how the order might be defined. If $\mathcal{X} \subseteq \mathbb{R}^d$, the usual componentwise order will be suitable in many applications, and we investigate it in Section 3.1. For covariates that are ordinal and admit a ranking in terms of importance, a lexicographic order may be suitable.

If groups of covariates are exchangeable, as in our case study on quantitative precipitation forecasts, other types of order relations need to be considered. In Sections 3.2 and 3.3, we study relations that are tailored to this setting, namely, the empirical stochastic order and empirical increasing convex order. Proofs are deferred to Supplementary Section S3.

3.1 | Componentwise order

Let $x = (x_1, \dots, x_d)$ and $x' = (x'_1, \dots, x'_d)$ denote elements of the covariate space \mathbb{R}^d . The most commonly used partial order in multivariate isotonic regression is the *componentwise order* defined by

$$x \leq x' \iff x_i \leq x'_i \text{ for } i = 1, \dots, d.$$

This order becomes weaker as the dimension d of the covariate space increases: If $\tilde{x} = (x_1, \dots, x_d, x_{d+1})$ and $\tilde{x}' = (x'_1, \dots, x'_d, x'_{d+1})$ then $x \leq x'$ is a necessary condition for $\tilde{x} \leq \tilde{x}'$. The following result is an immediate consequence of this observation and the structure of the optimization problem in Definition 1.

Proposition 1 *Let $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ have components $x_i = (x_{i1}, \dots, x_{id}) \in \mathbb{R}^d$ and $x_i^* = (x_{i1}, \dots, x_{id}, x_{i,d+1}) \in \mathbb{R}^{d+1}$ for $i = 1, \dots, n$, and let S be a proper scoring rule. Then if \mathbb{R}^d and \mathbb{R}^{d+1} are equipped with the componentwise partial order, and $\hat{\mathbf{F}}$ and $\hat{\mathbf{F}}^*$ denote S -based isotonic regressions of \mathbf{y} on \mathbf{x} and \mathbf{x}^* , respectively, it is true that*

$$\ell_S(\hat{\mathbf{F}}^*) \leq \ell_S(\hat{\mathbf{F}}).$$

In simple words, under the componentwise partial order, the inclusion of further covariates can only improve the in-sample fit. This behaviour resembles linear regression, where the addition of covariates can only improve the (unadjusted) R-square.

3.2 | Empirical stochastic order

We now define a relation that is based on stochastic dominance and invariant under permutation.

Definition 2 *Let $x = (x_1, \dots, x_d)$ and $x' = (x'_1, \dots, x'_d)$ denote elements of \mathbb{R}^d . Then x is smaller than or equal to x' in empirical stochastic order, for short $x \leq_{\text{st}} x'$, if the empirical distribution of x_1, \dots, x_d is smaller than the empirical distribution of x'_1, \dots, x'_d in the usual stochastic order.*

This relation is tailored to groups of exchangeable, real-valued covariates. The following result summarizes its properties and compares to the componentwise order, which we denote by \leq .

Proposition 2 Let $x = (x_1, \dots, x_d)$ and $x' = (x'_1, \dots, x'_d)$ denote elements of \mathbb{R}^d with order statistics $x_{(1)} \leq \dots \leq x_{(d)}$ and $x'_{(1)} \leq \dots \leq x'_{(d)}$.

- (a) The relation $x \leq_{\text{st}} x'$ is equivalent to $x_{(i)} \leq x'_{(i)}$ for $i=1, \dots, d$.
- (b) If $x \leq x'$ then $x \leq_{\text{st}} x'$.
- (c) If $x \leq_{\text{st}} x'$ and x and x' are comparable in the componentwise partial order, then $x \leq x'$.
- (d) If $x \leq_{\text{st}} x'$ and $x' \leq_{\text{st}} x$ then x and x' are permutations of each other. Consequently, the relation \leq_{st} defines a partial order on \mathbb{R}_+^d .

In a nutshell, the empirical stochastic order is equivalent to the componentwise order on the sorted elements, and this relation is weaker than the componentwise order. However, unlike the componentwise order, the empirical stochastic order does not degenerate as further covariates are added. To the contrary, empirical distributions of larger numbers of exchangeable variables become more informative and more easily comparable.

3.3 | Empirical increasing convex order

In applications, the empirical stochastic order might be too strong, in the sense that it does not generate sufficiently informative constraints. In this light, we now define a weaker partial order on \mathbb{R}_+^d , which also is based on a partial order for probability measures. Specifically, let X and X' be random variables with CDFs F and F' . Then F is smaller than F' in increasing convex order if $\mathbb{E}(\phi(X)) \leq \mathbb{E}(\phi(X'))$ for all increasing convex functions ϕ such that the expectations exist (Shaked & Shanthikumar, 2007, Section 4.A.1).

Definition 3 Let $x = (x_1, \dots, x_d)$ and $x' = (x'_1, \dots, x'_d)$ denote elements of \mathbb{R}^d . Then x is smaller than or equal to x' in empirical increasing convex order, for short $x \leq_{\text{icx}} x'$, if the empirical distribution of x_1, \dots, x_d is smaller than the empirical distribution of x'_1, \dots, x'_d in increasing convex order.

This notion provides another meaningful relation for groups of exchangeable covariates. The following result summarizes its properties and relates it to the empirical stochastic order.

Proposition 3 Let $x = (x_1, \dots, x_d)$ and $x' = (x'_1, \dots, x'_d)$ denote elements of \mathbb{R}^d with order statistics $x_{(1)} \leq \dots \leq x_{(d)}$ and $x'_{(1)} \leq \dots \leq x'_{(d)}$.

- (a) The relation $x \leq_{\text{icx}} x'$ is equivalent to

$$\sum_{i=j}^d x_{(i)} \leq \sum_{i=j}^d x'_{(i)} \text{ for } j = 1, \dots, d.$$

- (b) If $x \leq_{\text{st}} x'$ then $x \leq_{\text{icx}} x'$.
- (c) If $x \leq_{\text{icx}} x'$ then

$$\frac{1}{d} \sum_{i=1}^d x_i + \frac{d-1}{2(d+1)} g(x) \leq \frac{1}{d} \sum_{i=1}^d x'_i + \frac{d-1}{2(d+1)} g(x'),$$

where g is the Gini mean difference,




FIGURE 3 Regions of smaller, greater and incomparable elements in the positive quadrant of \mathbb{R}^2 , as compared to the point $(1, 3)$, for the (left) componentwise, (middle) empirical stochastic and (right) empirical increasing convex order. Coloured areas below (above) of $(1, 3)$ correspond to smaller (greater) elements, while blank areas contain elements incomparable to $(1, 3)$ in the given partial order [Colour figure can be viewed at wileyonlinelibrary.com]

$$g(x) = \frac{1}{d(d-1)} \sum_{i,j=1}^d |x_i - x_j|. \quad (22)$$

(d) If $x \leq_{\text{icx}} x'$ and $x' \leq_{\text{icx}} x$ then x and x' are permutations of each other. Consequently, the relation \leq_{icx} defines a partial order on \mathbb{R}_+^d .

Figure 3 illustrates the various types of relations for points in the positive quadrant of \mathbb{R}^2 . As reflected by the nested character of the regions, the componentwise order is stronger than the empirical stochastic order, which in turn is stronger than the empirical increasing convex order. The latter is equivalent to weak majorization as studied by Marshall et al. (2011). In the special case of vectors with non-negative entries, their Corollary C.5 implies that $x \in \mathbb{R}^d$ is dominated by $x' \in \mathbb{R}^d$ in empirical increasing convex order if, and only if, it lies in the convex hull of the points of the form $(\xi_1 x'_{\pi(1)}, \dots, \xi_d x'_{\pi(d)})$, where π is a permutation and $\xi_i \in \{0, 1\}$ for $i = 1, \dots, d$.

4 | SIMULATION STUDY

Since we view IDR primarily as a tool for prediction, we compare it to other distributional regression methods in terms of predictive performance on continuous and discrete, univariate simulation examples, as measured by the CRPS. However, as noted below and formalized in Supplementary Section S4, the CRPS links asymptotically to L_2 estimation error, so under large validation samples, prediction and estimation are assessed simultaneously. A detailed comparative study on mixed discrete-continuous data with a multivariate covariate vector is given in the case study in the next section.

Here, our simulation scenarios build on the illustrating example in the introduction. Specifically, we draw a covariate $X \sim \text{Unif}(0, 10)$ and then

$$Y_1 | X \sim \text{Gamma}(\text{shape} = \sqrt{X}, \text{scale} = \min\{\max\{X, 1\}, 6\}), \quad (23)$$

$$Y_2 | X = Y_1 | X + 10 \cdot \mathbb{1}\{X \geq 5\}, \quad (24)$$

TABLE 1 Mean CRPS in smooth (23), discontinuous (24), non-isotonic (25), and discrete (26) simulation scenarios with training sets of size n

n	500	1000	2000	4000	500	1000	2000	4000
Smooth (23)								
NP	3.561	3.542	3.532	3.525	Discontinuous (24)			
SQR	3.571	3.543	3.530	3.524	3.614	3.582	3.562	3.549
TRAM	3.560	3.543	3.535	3.531	3.647	3.619	3.606	3.600
QRF	3.589	3.561	3.555	3.553	3.642	3.625	3.616	3.612
IDR	3.604	3.568	3.548	3.535	3.614	3.576	3.561	3.556
IDR _{sbg}	3.595	3.561	3.543	3.532	3.628	3.581	3.555	3.540
Ideal	3.516	3.516	3.516	3.516	3.620	3.577	3.551	3.537
Non-isotonic (25)								
Discrete (26)								
NP	3.564	3.544	3.534	3.527	3.516	3.516	3.516	3.516
SQR	3.574	3.546	3.533	3.527	1.136	1.131	1.128	1.126
TRAM	3.566	3.549	3.543	3.539	1.129	1.121	1.116	1.114
QRF	3.587	3.560	3.555	3.553	1.115	1.110	1.107	1.106
IDR	3.605	3.569	3.549	3.536	1.121	1.113	1.112	1.112
IDR _{sbg}	3.597	3.564	3.545	3.534	1.130	1.119	1.113	1.109
Ideal	3.516	3.516	3.516	3.516	1.128	1.118	1.112	1.109
					1.104	1.104	1.104	1.104

$$Y_3 | X = Y_1 | X - 2 \cdot \mathbb{1}\{X \geq 7\}, \quad (25)$$

$$Y_4 | X \sim \text{Poisson}(\lambda = \min\{\max\{X, 1\}, 6\}). \quad (26)$$

Under each scenario, we generate 500 training sets of size $n = 500, 1000, 2000$ and 4000 each, fit distributional regression models, and validate on a test set of size $m = 5000$. For comparison with IDR, we use a non-parametric kernel (or nearest neighbour) smoothing technique (NP; Li & Racine, 2008), semiparametric quantile regression with monotone rearrangement (SQR; Chernozhukov et al., 2010; Koenker, 2005), conditional transformation models (TRAM; Hothorn et al., 2014) and distributional or quantile random forests (QRF; Athey et al., 2019; Meinshausen, 2006). These methods have been chosen as they are not subject to restrictive assumptions on the distribution of the response variable and have well-established and well-documented implementations in the statistical programming environment R (R Core Team, 2020). We also include the ideal forecast, that is, the true conditional distribution of the response given the covariate, in the comparison.

Implementation details for the various methods are given in Table S1 in Supplementary Section S5. Here we only note that QRF uses the `grf` package (Tibshirani et al., 2020) with a splitting rule that is tailored to quantiles (Athey et al., 2019). We see that, unlike IDR, its competitors rely on manual intervention and tuning. For example, QRFs perform poorly under the default value of 5 for the tuning parameter `min.node.size`, which we have raised to 40. Further improvement may arise when tuning parameters, such as honesty fraction and node size, are judiciously adjusted to the specific scenario and training sample size at hand. In contrast, IDR is entirely free of implementation decisions, except for the subbagging variant, IDR_{sbg} , where we average predictions based on estimates on 100 subsamples of size $n/2$ each.

Table 1 shows the mean CRPS for the different methods and simulation scenarios. Scenario (23) is the same as in the introduction and illustrated in Figure 1. It has a smooth covariate–response relationship, and NP, SQR, and even the misspecified TRAM technique, which are tailored to this type of setting, outperform QRF and IDR. However, the assumption of continuity in the response is crucial, as the results under the discontinuous scenario (24) demonstrate, where IDR and IDR_{sbg} perform best. In the non-isotonic scenario (25), IDR and IDR_{sbg} retain acceptable performance, even though the key assumption is violated. Not surprisingly, SQR faces challenges in the Poisson scenario (26), where the conditional quantile functions are piecewise constant, and IDR is outperformed only by TRAM. Throughout, the simplistic subbagging variant of IDR has slightly lower mean CRPS than the default variant that is estimated on the full training set, and it would be interesting to explore the relation to the super-efficiency phenomenon described by Banerjee et al. (2019).

These results lend support to our belief that IDR can serve as a universal benchmark in probabilistic forecasting and distributional regression problems. For sufficiently large training samples, IDR offers competitive performance under any type of linearly ordered outcome, without reliance on tuning parameters or other implementation choices, except when subsampling is employed.

TABLE 2 Meteorological stations at airports, with International Air Transport Association (IATA) airport code, World Meteorological Organization (WMO) station ID, and data availability in days (years)

	IATA code	WMO ID	Data availability
Brussels, Belgium	BRU	06449	3406 (9.3)
Frankfurt, Germany	FRA	10637	3617 (9.9)
London, UK	LHR	03772	2256 (6.2)
Zurich, Switzerland	ZRH	06670	3241 (8.9)

5 | CASE STUDY: PROBABILISTIC QUANTITATIVE PRECIPITATION FORECASTS

The past decades have witnessed tremendous progress in the science and practice of weather prediction (Bauer et al., 2015). Arguably, the most radical innovation consists in the operational implementation of ensemble systems and an accompanying culture change from point forecasts to distributional forecasts (Leutbecher & Palmer, 2008). An ensemble system comprises multiple runs of numerical weather prediction (NWP) models, where the runs or members differ from each other in initial conditions and numerical-physical representations of atmospheric processes.

Ideally, one would like to interpret an ensemble forecast as a random sample from the conditional distribution of future states of the atmosphere. However, this is rarely advisable in practice, as ensemble forecasts are subject to biases and dispersion errors, thereby calling for some form of statistical post-processing (Gneiting & Raftery, 2005; Vannitsem et al., 2018). This is typically done by fitting a distributional regression model, with the weather variable of interest being the response variable, and the members of the forecast ensemble constituting the covariates, and applying this model to future NWP output, to obtain conditional predictive distributions for future weather quantities. State of the art techniques include Bayesian Model Averaging (BMA; Raftery et al., 2005; Sloughter et al., 2007), Ensemble Model Output Statistics (EMOS; Gneiting et al., 2005; Scheuerer, 2014), and Heteroscedastic Censored Logistic Regression (HCLR; Messner et al., 2014).

In this case study, we apply IDR to the statistical post-processing of ensemble forecasts of accumulated precipitation, a variable that is notoriously difficult to handle, due to its mixed discrete-continuous character, which requires both a point mass at zero and a right skewed continuous component on the positive half-axis. As competitors to IDR, we implement the BMA technique of Sloughter et al. (2007), the EMOS method of Scheuerer (2014), and HCLR (Messner et al., 2014), which are widely used parametric approaches that have been developed specifically for the purposes of probabilistic quantitative precipitation forecasting. In contrast, IDR is a generic technique and fully automatic, once the partial order on the covariate space has been specified.

5.1 | Data

The data in our case study comprise forecasts and observations of 24 h accumulated precipitation from 6 January 2007 to 1 January 2017 at meteorological stations on airports in London, Brussels, Zurich and Frankfurt. As detailed in Table 2, data availability differs, and we remove days with missing entries station by station, so that all types of forecasts for a given station are trained

and evaluated on the same data. Both forecasts and observations refer to the 24 h period from 6:00 UTC to 6:00 UTC on the following day. The observations are in the unit of millimetre and constitute the response variable in distributional regression. They are typically, but not always, reported in integer multiples of a millimetre.

As covariates, we use the 52 members of the leading NWP ensemble operated by the European Centre for Medium-Range Weather Forecasts (ECMWF; Buizza et al., 2005; Molteni et al., 1996). The ECMWF ensemble system comprises a high-resolution member (x_{HRES}), a control member at lower resolution (x_{CTR}) and 50 perturbed members (x_1, \dots, x_{50}) at the same lower resolution but with perturbed initial conditions, and the perturbed members can be considered exchangeable (Leutbecher, 2019). To summarize, the covariate vector in distributional regression is

$$x = (x_1, \dots, x_{50}, x_{\text{CTR}}, x_{\text{HRES}}) = (x_{\text{PTB}}, x_{\text{CTR}}, x_{\text{HRES}}) \in \mathbb{R}^{52}, \quad (27)$$

where $x_{\text{PTB}} = (x_1, \dots, x_{50}) \in \mathbb{R}^{50}$. At each station, we use the forecasts for the corresponding latitude-longitude gridbox of size 0.25×0.25 degrees, and we consider prediction horizons of 1 to 5 days. For example, the two day forecast is initialized at 00:00 Universal Coordinated Time (UTC) and issued for 24 h accumulated precipitation from 06:00 UTC on the next calendar day to 06:00 UTC on the day after. ECMWF forecast data are available through the ECMWF Meteorological Archival and Retrieval System (MARS; <https://www.ecmwf.int/en/forecasts>) and via the TIGGE system (Bougeault et al., 2010; Swinbank et al., 2016). Observation data are provided by NOAA's Integrated Surface Database (ISD; <https://www.ncdc.noaa.gov/isd>).

Statistical post-processing is both a calibration and a downscaling problem: Forecasts and observations are at different spatial scales, whence the unprocessed forecasts are subject to representativeness error (Wilks, 2019, Chapter 8.9). Indeed, if we interpret the predictive distribution from the raw ensemble (27) as the empirical distribution of all 52 members—a customary approach, which we adopt hereinafter—there is a strong bias in probability of precipitation forecasts: Days with exactly zero precipitation are predicted much less often at the NWP model grid box scale than they occur at the point scale of the observations.

5.2 | BMA, EMOS and HCLR

Before describing our IDR implementation, we review its leading competitors, namely, state of the art parametric distributional regression approaches that have been developed specifically for accumulated precipitation.

Techniques of ensemble model output statistics (EMOS; Gneiting et al., 2005) type can be interpreted as parametric instances of generalized additive models for location, scale and shape (GAMLSS; Rigby & Stasinopoulos, 2005). The specific variant of Scheuerer (2014) which we use here is based on the three-parameter family of left-censored generalized extreme value (GEV) distributions. The left-censoring generates a point mass at zero, corresponding to no precipitation, and the shape parameter allows for flexible skewness on the positive half-axis, associated with rain, hail or snow accumulations. The GEV location parameter is modelled as a linear function of $x_{\text{HRES}}, x_{\text{CTR}}, m_{\text{PTB}} = \frac{1}{50} \sum_{i=1}^{50} x_i$

$$p_{\text{ZERO}} = \frac{1}{52} \left(\mathbb{1}\{x_{\text{HRES}} = 0\} + \mathbb{1}\{x_{\text{CTR}} = 0\} + \sum_{i=1}^{50} \mathbb{1}\{x_i = 0\} \right),$$

and the GEV scale parameter is linear in the Gini mean difference (22) of the 52 individual forecasts in the covariate vector (27). While all parameters are estimated by minimizing the in-sample CRPS, the GEV shape parameter does not link to the covariates.

The general idea of the Bayesian model averaging (BMA; Raftery et al., 2005) approach is to employ a mixture distribution, where each mixture component is parametric and associated with an individual ensemble member forecast, with mixture weights that reflect the member's skill. Here we use the BMA implementation of Sloughter et al. (2007) for accumulated precipitation in a variant that is based on $x_{\text{HRES}}, x_{\text{CTR}}, m_{\text{PTB}} = \frac{1}{50} \sum_{i=1}^{50} x_i$ which we found to achieve more stable estimates and superior predictive scores than variants based on all members, as proposed by Fraley et al. (2010) in settings with smaller groups of exchangeable members. Hence, our BMA predictive CDF is of the form

$$F_x(y) = w_{\text{HRES}} G(y|x_{\text{HRES}}) + w_{\text{CTR}} G(y|x_{\text{CTR}}) + w_{\text{PTB}} G(y|x_{\text{PTB}})$$

for $y \in \mathbb{R}$, where the component CDFs $G(y|\cdot)$ are parametric, and the weights $w_{\text{HRES}}, w_{\text{CTR}}$ and w_{PTB} are non-negative and sum to one. Specifically, $G(y|x_{\text{HRES}})$ models the logit of the point mass at zero as a linear function of $\sqrt[3]{x_{\text{HRES}}}$ and $p_{\text{HRES}} = \mathbb{1}\{x_{\text{HRES}} = 0\}$, and the distribution for positive accumulations as a gamma density with mean and variance being linear in $\sqrt[3]{x_{\text{HRES}}}$ and x_{HRES} , respectively, and analogously for $G(y|x_{\text{CTR}})$ and $G(y|x_{\text{PTB}})$. Estimation relies on a two-step procedure, where the (component specific) logit and mean models are fitted first, followed by maximum likelihood estimation of the weight parameters and the (joint) variance model via the EM algorithm (Sloughter et al., 2007).

Heteroscedastic censored logistic regression (Messner et al., 2014) originates from the observation that conditional CDFs can be estimated by dichotomizing the random variable of interest at given thresholds and estimating the probability of threshold exceedance via logistic regression. The HCLR model used here assumes that square-root transformed precipitation follows a logistic distribution censored at zero, with location parameter linear in $\sqrt{x_{\text{HRES}}}, \sqrt{x_{\text{CTR}}}$ and the mean of the square-root transformed perturbed forecasts, and a scale parameter linear in the standard deviation of the square-root transformed perturbed forecasts. Like EMOS, HCLR can be interpreted within the GAMLSS framework of Rigby & Stasinopoulos (2005).

Code for BMA, EMOS and HCLR is available within the `ensembleBMA`, `ensembleMOS` and `crch` packages in R (Messner, 2018). Unless noted differently, we use default options in implementation decisions.

5.3 | Choice of partial order for IDR

IDR applies readily in this setting, without any need for adaptations due to the mixed-discrete continuous character of precipitation accumulation, nor requiring data transformations or other types of implementation decisions. However, the partial order on the elements (27) of the covariate space $\mathcal{X} = \mathbb{R}^{52}$, or on a suitable derived space, needs to be selected thoughtfully, considering that the perturbed members x_1, \dots, x_{50} are exchangeable.

In the sequel, we apply IDR in three variants. Our first implementation is based on $x_{\text{HRES}}, x_{\text{CTR}}$ and $m_{\text{PTB}} = \frac{1}{50} \sum_{i=1}^{50} x_i$ along with the componentwise order on \mathbb{R}^3 , in that

$$x \leq x' \iff m_{\text{PTB}} \leq m'_{\text{PTB}}, x_{\text{CTR}} \leq x'_{\text{CTR}}, x_{\text{HRES}} \leq x'_{\text{HRES}}. \quad (28)$$

The second implementation uses the same variables and partial order, but combined with a simple subbagging approach: Before applying IDR, the training data is split into the two disjoint subsamples of training observations with odd and even indices, and we average the predictions based on these two subsamples.

Our third implementation combines the empirical increasing convex order for x_{PTB} with the usual total order on \mathbb{R} for x_{HRES} , whence

$$x \leq x' \iff x_{\text{PTB}} \leq_{\text{icx}} x'_{\text{PTB}}, x_{\text{HRES}} \leq x'_{\text{HRES}}. \quad (29)$$

Henceforth, we refer to the three implementations based on the partial orders in (28) and (29) as IDR_{cw} , IDR_{sbg} , and IDR_{icx} . With reference to the discussion preceding Theorem 1, the relations (28) and (29) define preorders on \mathbb{R}^{52} and partial orders on \mathbb{R}^3 and $\mathbb{R}_1^{50} \times \mathbb{R}$, respectively.

We have experimented with other options as well, for example, by incorporating the maximum $\max_{i=1, \dots, 50} x_i$ of the perturbed members in the componentwise order in (28), with the motivation that the maximum might serve as a proxy for the spread of the ensemble, or by using the empirical stochastic order \leq_{st} in lieu of the empirical increasing convex order \leq_{icx} in (29). IDR is robust to changes of this type, and the predictive performance remains stable, provided that the partial order honors the key substantive insights, in that the perturbed members x_1, \dots, x_{50} are exchangeable, while x_{HRES} , due to its higher native resolution, is able to capture local information that is not contained in x_{PTB} nor x_{CTR} . Hence, x_{HRES} ought to play a pivotal role in the partial order.

5.4 | Selection of training periods

The selection of the training period is a crucial step in the statistical post-processing of NWP output. Most post-processing methods, including the ones used in this analysis, assume that there is a stationary relationship between the forecasts and the observations. As Hamill (2018) points out, this assumption is hardly ever satisfied in practice: NWP models are updated, instruments at observation stations get replaced, and forecast biases may vary seasonally. These problems are exacerbated by the fact that quantitative precipitation forecasts require large training datasets in order to include sufficient numbers of days with non-zero precipitation and extreme precipitation events.

For BMA and EMOS, a training period over a rolling window of the latest available 720 days at the time of forecasting is (close to) optimal at all stations. This resembles choices made by Scheuerer & Hamill (2015) who used a training sample of about 900 past instances. Scheuerer (2014) took shorter temporal windows, but merged instances from nearby stations into the training sets, which is not possible here. In general, it would be preferable to select training data seasonally (e.g. data from the same month), but in our case the positive effect of using seasonal training data does not outweigh the negative effect of a smaller sample size.

As a non-parametric technique, IDR requires larger sets of training data than BMA or EMOS. As training data for IDR, we used all data available at the time of forecasting, which is about 2500 to 3000 days for the stations Frankfurt, Brussels and Zurich, and 1500 days for London Heathrow. The same training periods are also used for HCLR, where no positive effect of shorter, rolling training periods has been observed (Messner et al., 2014).




FIGURE 4 Distributional forecasts for accumulated precipitation at Brussels, valid 16 December 2015 at a prediction horizon of 2 days. (a) BMA, EMOS, IDR_{cw} , IDR_{sbg} and IDR_{icx} predictive CDFs. The vertical line represents the observation. (b) IDR_{cw} CDF along with the hard and soft constraints in (20) as induced by the order relation (28). The thin lines show the IDR_{cw} CDFs at direct predecessors and successors [Colour figure can be viewed at wileyonlinelibrary.com]

For evaluation, we use the years 2015 and 2016 (and 01 January 2017) for all post-processing methods and the raw ensemble. This test dataset consists of roughly 700 instances for each station and lead time.

5.5 | Results

Before comparing the BMA, EMOS, IDR_{cw} , IDR_{sbg} and IDR_{icx} techniques in terms of out-of-sample predictive performance over the test period, we exemplify them in Figure 4, where we show predictive CDFs for accumulated precipitation at Brussels on 16 December 2015, at a prediction horizon of 2 days. In panel (a), the marks at the bottom correspond to x_{HRES} , x_{CTR} , the perturbed members x_1, \dots, x_{50} and their mean m_{PTB} . The observation at 4 mm is indicated by the vertical line. Under all four techniques, the point mass at zero, which represents the probability of no precipitation, is vanishingly small. While the BMA, EMOS and HCLR CDFs are smooth and supported on the positive half-axis, the IDR_{cw} , IDR_{sbg} and IDR_{icx} CDFs are piecewise constant with jump points at observed values in the training period. Panel (b) illustrates the hard




FIGURE 5 Mean CRPS over the test period for raw and post-processed ensemble forecasts of 24 h accumulated precipitation at prediction horizons of 1, 2, 3, 4 and 5 days. The lowest mean score for a given lead time and station is indicated in green [Colour figure can be viewed at wileyonlinelibrary.com]

and soft constraints on the IDR_{cw} CDF that arise from (20) under the order relation (28), with the thinner lines representing the IDR_{cw} CDFs of direct successors and predecessors. In this example, the constraints are mostly hard, except for threshold values between 4 and 11 mm.

We now use the mean CRPS over the test period as an overall measure of out-of-sample predictive performance. Figure 5 shows the CRPS of the raw and post-processed forecasts for all stations and lead times, with the raw forecast denoted as ENS. While HCLR performs best in terms of the CRPS, the IDR variants show scores of a similar magnitude and outperform BMA in many instances. Figure S1 in Supplementary Section S5 shows the difference of the empirical cumulative distribution function (ECDF) of the PIT defined at (8) to the bisector for the distributional forecasts. All three IDR variants show a PIT-distribution close to uniform, and so do BMA, EMOS and HCLR, as opposed to the raw ensemble, which is underdispersed.

In Figure 6, we evaluate probability of precipitation forecasts by means of the Brier score (Gneiting & Raftery, 2007), and Figure S2 in Supplementary Section S5 shows reliability diagrams (Dimitriadis et al., 2021; Wilks, 2019). As opposed to the raw ensemble forecast, all distributional regression methods yield reliable probability forecasts. BMA, IDR_{cw} , IDR_{sbg} and IDR_{icx} separate the estimation of the point mass at zero, and of the distribution for positive accumulations, and the four methods perform ahead of EMOS. HCLR is outperformed by BMA and the IDR variants at lead times of one or two days, but achieves a lower Brier score at the longest lead time of 5 days.

Interestingly, IDR tends to outperform EMOS and HCLR for probability of precipitation forecasts, but not for precipitation accumulations. We attribute this to the fact that parametric techniques are capable of extrapolating beyond the range of the training responses, whereas IDR is not: The highest precipitation amount judged feasible by IDR equals the largest observation in the training set. Furthermore, unlike EMOS and HCLR, IDR does not use information about the




FIGURE 6 Mean Brier score over the test period for probability of precipitation forecasts at prediction horizons of 1, 2, 3, 4 and 5 days. The lowest mean score for a given lead time and station is indicated in green [Colour figure can be viewed at wileyonlinelibrary.com]

spread of the raw ensemble, which is inconsequential for probability of precipitation forecasts, but may impede distributional forecasts of precipitation accumulations.

In all comparisons, the forecast performance of IDR_{cw} and IDR_{sbg} is similar. However, in our implementation, the simple subagging method used in IDR_{sbg} reduced the computation time by up to one half.

To summarize, our results underscore the suitability of IDR as a benchmark technique in probabilistic forecasting problems. Despite being generic as well as fully automated, IDR is remarkably competitive relative to state of the art techniques that have been developed specifically for the purpose. In fact, in a wide range of applied problems that lack sophisticated, custom-made distributional regression solutions, IDR might well serve as a ready-to-use, top-performing method of choice.

6 | DISCUSSION

Stigler (1975) gives a lucid historical account of the 19th century transition from point estimation to distribution estimation. In regression analysis, we may be witnessing what future generations might refer to as the transition from conditional mean estimation to conditional distribution estimation, accompanied by a simultaneous transition from point forecasts to distributional forecasts (Gneiting & Katzfuss, 2014).

IDR is a non-parametric technique for estimating conditional distributions that takes advantage of partial order relations within the covariate space. It can be viewed as a far-reaching generalization of pool adjacent violators (PAV) algorithm based classical approaches to isotonic (non-distributional) regression, is entirely generic and fully automated, and provides for a unified

treatment of continuous, discrete and mixed discrete-continuous real-valued response variables. Code for the implementation of IDR within R (R Core Team, 2020) and Python (<https://www.python.org/>) is available via the `isodistrreg` package at CRAN (<https://CRAN.R-project.org/package=isodistrreg>) and on github (<https://github.com/AlexanderHenzi/isodistrreg>; <https://github.com/evwalz/isodisreg>), with user-friendly functions for partial orders, estimation, prediction and evaluation.

IDR relies on information supplied by order constraints, and the choice of the partial order on the covariate space is a critical decision prior to the analysis. Only variables that contribute to the partial order need to be retained, and the order constraints serve to regularize the IDR solution. Weak orders lead to increased numbers of comparable pairs of training instances and predictive distributions that are more regular. The choice of the partial order is typically guided and informed by substantive expertise, as illustrated in our case study, and it is a challenge for future research to investigate whether the selection of the partial order could be automated. Given that IDR gains information through order constraints, it is a valid concern whether it is robust under misspecifications of the partial order. There is evidence that this is indeed the case: IDR has guaranteed in-sample threshold calibration (Theorem 2) and therefore satisfies a minimal requirement for reliable probabilistic forecasts under any (even misspecified) partial order. Moreover, El Barmi & Mukerjee (2005, Theorem 7) show that in the special case of a discrete, totally ordered covariate, isotonic regression asymptotically has smaller estimation error than non-isotonic alternatives even under mild violations of the monotonicity assumptions, akin to the performance of IDR in the non-isotonic setting (25) in our simulation study.

Unlike other methods for distributional regression, which require implementation decisions, such as the specification of parametric distributions, link functions, estimation procedures and convergence criteria, to be undertaken by users, IDR is fully automatic once the partial order and the training set have been identified. In this light, we recommend that IDR be used as a benchmark technique in distributional regression and probabilistic forecasting problems. With both computational efficiency and the avoidance of overfitting in mind, IDR can be combined with subsample aggregation (subagging) in the spirit of random forests. In our case study on quantitative precipitation forecasts, we used simplistic ad hoc choices for the size and number of subsamples. Future research on computationally efficient algorithmic implementations of IDR as well as optimal and automated choices of subsampling settings is highly desirable.

A limitation of IDR in its present form is that we only consider the usual stochastic order on the space \mathcal{P} of the conditional distributions. Hence, IDR is unable to distinguish situations where the conditional distributions agree in location but differ in spread, shape or other regards. This restriction is of limited concern for response variables such as precipitation accumulation or income, which are bounded below and right skewed, but may impact the application of IDR to variables with symmetric distributions. In this light, we encourage future work on ramifications of IDR, in which \mathcal{P} is equipped with partial orders other than the stochastic order, including but not limited to the likelihood ratio order (Mösching & Dümbgen, 2020a). Similarly, the ‘spiking’ problem of traditional isotonic regression, which refers to unwarranted jumps of estimates at boundaries, arguably did not have adverse effects in our simulation and case studies. However, it might be of concern in other applications, where remedies of the type proposed by Wu et al. (2015) might yield improvement and warrant study.

Another promising direction for further research is generalizations of IDR to multivariate response variables. In weather prediction, this would allow simultaneous post-processing of forecasts for several variables, and an open question is for suitable notions of multivariate stochastic dominance that allow efficient estimation in such settings.

ACKNOWLEDGEMENTS

The authors are grateful to the editor, the associate editor, four anonymous referees, Sebastian Arnold, Lutz Dümbgen, Ed George, Stephan Hemri, Alexander Jordan, Kristof Kraus, Alexandre Mösching, Anja Mühlemann, Aaditya Ramdas, Nicholas Reich, Benedikt Schulz, Peter Vogel, Jonas Wallin and Eva-Maria Walz for providing code, assistance with the data handling, comments and discussion. Tilmann Gneiting acknowledges support by the Klaus Tschira Foundation, by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 257899354 – TRR 165, and via the Fellowship programme operated by the European Centre for Medium-Range Weather Forecasts (ECMWF). Alexander Henzi and Johanna F. Ziegel gratefully acknowledge financial support from the Swiss National Science Foundation. Calculations were mostly performed on UBELIX (<http://www.id.unibe.ch/hpc>), the HPC cluster at the University of Bern.

ORCID

Johanna F. Ziegel <http://orcid.org/0000-0002-5916-9746>

Tilmann Gneiting <http://orcid.org/0000-0001-9397-3271>

REFERENCES

- Allen, M.S. & Eckel, F.A. (2012) Value from ambiguity in ensemble forecasts. *Weather and Forecasting*, 27, 70–84.
- Athey, S., Tibshirani, J. & Wager, S. (2019) Generalized random forests. *Annals of Statistics*, 37, 1148–1178.
- Ayer, M., Brunk, H.D., Ewing, G.M., Reid, W.T. & Silverman, E. (1955) An empirical distribution function for sampling with incomplete information. *Annals of Mathematical Statistics*, 26, 641–647.
- Banerjee, M., Durot, C. & Sen, B. (2019) Divide and conquer in nonstandard problems and the super efficiency phenomenon. *Annals of Statistics*, 47, 720–757.
- Barlow, R.E., Bartholomew, D.J., Bremner, J.M. & Brunk, H.D. (1972) *Statistical inference under order restrictions*. Hoboken: John Wiley & Sons.
- Bartholomew, D.J. (1959a) A test of homogeneity for ordered alternatives. *Biometrika*, 46, 36–48.
- Bartholomew, D.J. (1959b) A test of homogeneity for ordered alternatives II. *Biometrika*, 46, 329–335.
- Basel Committee on Banking Supervision. (2016) Standard: minimum capital requirements for market risk. Bank for International Settlements. Available from: <http://www.bis.org/bcbs/publ/d352.pdf>.
- Bauer, P., Thorpe, A. & Brunet, G. (2015) The quiet revolution of numerical weather prediction. *Nature*, 525, 47–55.
- Ben Bouallègue, Z., Haiden, T. & Richardson, D.S. (2018) The diagonal score: definition, properties, and interpretations. *Quarterly Journal of the Royal Meteorological Society*, 144, 1463–1473.
- Bougeault, P., Toth, Z., Bishop, C., Brown, B., Burridge, D., Chen, D.H. et al. (2010) The THORPEX interactive grand global ensemble. *Bulletin of the American Meteorological Society*, 91, 1059–1072.
- Breiman, L. (1996) Bagging predictors. *Machine Learning*, 24, 123–140.
- Breiman, L. (2001) Random forests. *Machine Learning*, 45, 5–32.
- Breiman, L., Friedman, J.H., Olshen, R.A. & Stone, C.J. (1984) *Classification and regression trees*. USA: Wadsworth.
- Brightwell, G. (1992) Random k-dimensional orders: Width and number of linear extensions. *Order*, 9, 333–342.
- Brunk, H.B. (1955) Maximum likelihood estimates of monotone parameters. *Annals of Mathematical Statistics*, 26, 607–616.
- Bühlmann, P. & Yu, B. (2002) Analyzing bagging. *Annals of Statistics*, 30, 927–961.
- Buizza, R., Houtekamer, P. L., Pellerin, G., Toth, Z., Zhu, Y. & Wei, M. (2005) A comparison of the ECMWF, MSC, and NCEP global ensemble prediction systems. *Monthly Weather Review*, 133, 1076–1097.
- Buja, A. & Stützel, W. (2006) Observations on bagging. *Statistica Sinica*, 16, 323–351.
- Casady, R.J. & Cryer, J.D. (1976) Monotone percentile regression. *Annals of Statistics*, 4, 532–541.
- Chernozhukov, V., Fernández-Val, I. & Galichon, A. (2010) Quantile and probability curves without crossing. *Econometrica*, 78, 1093–1125.
- Davidov, O. & Iliopoulos, G. (2012) Estimating a distribution function subject to a stochastic ordering restriction: a comparative study. *Journal of Nonparametric Statistics*, 24, 923–933.

- Dawid, A.P. (1984) Statistical theory: the prequential approach. *Journal of the Royal Statistical Society Series A*, 147, 278–290.
- Dette, H., Neumeyer, N. & Pilz, K.F. (2006) A simple nonparametric estimator of a strictly monotone regression function. *Bernoulli*, 12, 469–490.
- Diebold, F. X., Gunther, T.A. & Tay, A.S. (1998) Evaluating density forecasts with applications to financial risk management. *International Economic Review*, 39, 863–883.
- Dimitriadis, T., Gneiting, T. & Jordan, A.I. (2021) Stable reliability diagrams for probabilistic classifiers. *Proceedings of the National Academy of Sciences*, 118, e2016191118.
- Dunson, D.B., Pillai, N. & Park, J.-H. (2007) Bayesian density regression. *Journal of the Royal Statistical Society Series B*, 69, 163–183.
- van Eeden, C. (1958) Testing and estimating ordered parameters of probability distributions. Ph. D. thesis, University of Amsterdam.
- Ehm, W., Gneiting, T., Jordan, A. & Krüger, F. (2016) Of quantiles and expectiles: Consistent scoring functions, Choquet representations, and forecast rankings (with discussion). *Journal of the Royal Statistical Society Series B*, 78, 505–562.
- El Barmi, H. & Mukerjee, H. (2005) Inferences under a stochastic ordering constraint. *Journal of the American Statistical Association*, 100, 252–261.
- Ellsberg, D. (1961) Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics* 75, 643–669.
- Fraley, C., Raftery, A.E. & Gneiting, T. (2010) Calibrating multimodel forecast ensembles with exchangeable and missing members using Bayesian model averaging. *Monthly Weather Review*, 138, 190–202.
- Gasthaus, J., Benidis, K., Wang, Y., Rangapuram, S.S., Salinas, D., Flunkert, V. et al. (2019) Probabilistic forecasting with spline quantile function RNNs. *Proceedings of Machine Learning Research*, 89, 1901–1910.
- Gneiting, T. (2011) Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106, 746–762.
- Gneiting, T. & Katzfuss, M. (2014) Probabilistic forecasting. *Annual Review of Statistics and its Application*, 1, 125–151.
- Gneiting, T. & Raftery, A.E. (2005) Weather forecasting with ensemble methods. *Science*, 310, 248–249.
- Gneiting, T. & Raftery, A.E. (2007) Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 359–378.
- Gneiting, T. & Ranjan, R. (2011) Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, 29, 411–422.
- Gneiting, T. & Ranjan, R. (2013) Combining predictive distributions. *Electronic Journal of Statistics*, 7, 1747–1782.
- Gneiting, T., Raftery, A.E., Westveld III, A.H. & Goldman, T. (2005) Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review*, 133, 1098–1118.
- Gneiting, T., Balabdaoui, F. & Raftery, A.E. (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society Series B*, 69, 243–268.
- Groeneboom, P. & Jongbloed, G. (2014) *Nonparametric estimation under shape constraints*. Cambridge: Cambridge University Press.
- Guntuboyina, A. & Sen, B. (2018) Nonparametric shape-restricted regression. *Statistical Science*, 33, 563–594.
- Hall, P., Wolff, R.C.L. & Yao, Q. (1999) Methods for estimating a conditional distribution function. *Journal of the American Statistical Association*, 94, 154–163.
- Hamill, T.M. (2018) Practical aspects of statistical postprocessing. In: Vannitsem, S., Wilks, D.S. & Messner, J. (Eds.) *Statistical Postprocessing of Ensemble Forecasts*. Amsterdam: Elsevier, pp. 187–217.
- Henzi, A., Mösching, A. & Dümbgen, L. (2020) Accelerating the pool-adjacent-violators algorithm for isotonic distributional regression. Preprint, arxiv.org/abs/2006.05527.
- Hersbach, H. (2000) Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather and Forecasting*, 15, 559–570.
- Hogg, R.V. (1965) On models and hypotheses with restricted alternatives. *Journal of the American Statistical Association*, 60, 1153–1162.
- Hothorn, T., Lausen, B., Benner, A. & Radespiel-Tröger, M. (2004) Bagging survival trees. *Statistics in Medicine*, 23, 77–91.
- Hothorn, T., Kneib, T. & Bühlmann, P. (2014) Conditional transformation models. *Journal of the Royal Statistical Society Series B*, 76, 3–27.

- Jordan, A.I., Mühlemann, A. & Ziegel, J.F. (2021) Optimal solutions to the isotonic regression problem. *Annals of the Institute of Statistical Mathematics*. To appear. Preprint, arxiv.org/abs/1904.04761.
- Klein, N., Kneib, T., Lang, S. & Sohn, A. (2015) Bayesian structured additive distributional forecasting with an application to regional income inequality in Germany. *Annals of Applied Statistics*, 9, 1024–1052.
- Koenker, R. (2005) *Quantile regression*. Cambridge: Cambridge University Press.
- Laio, F. & Tamea, S. (2007) Verification tools for probabilistic forecasts of continuous hydrological variables. *Hydrology and Earth System Sciences*, 11, 1267–1277.
- de Leeuw, J., Hornik, K. & Mair, P. (2009) Isotone optimization in R: pool-adjacent-violators algorithm (PAVA) and active set methods. *Journal of Statistical Software*, 32, 1–19.
- Leutbecher, M. (2019) Ensemble size: How suboptimal is less than infinity? *Quarterly Journal of the Royal Meteorological Society*, 145, 107–128.
- Leutbecher, M. & Palmer, T.N. (2008) Ensemble forecasting. *Journal of Computational Physics*, 227, 3515–3539.
- Li, Q. & Racine, J. (2008) Nonparametric estimation of conditional CDF and quantile functions with mixed categorical and continuous data. *Journal of Business & Economic Statistics*, 26, 423–434.
- Mammen, E. (1991) Estimating a smooth monotone regression function. *Annals of Statistics*, 19, 724–740.
- Marshall, A.W., Olkin, I. & Arnold, B.C. (2011) *Inequalities: theory of majorization and its applications*, 2nd edn. Berlin: Springer.
- Matheson, J.E. & Winkler, R.L. (1976) Scoring rules for continuous probability distributions. *Management Science*, 22, 1087–1096.
- Meinshausen, N. (2006) Quantile regression forests. *Journal of Machine Learning Research*, 7, 983–999.
- Messner, J.W. (2018) Ensemble postprocessing with R. In: Vannitsem, S., Wilks, D.S. & Messner, J. (Eds.) *Statistical Postprocessing of Ensemble Forecasts*. Amsterdam: Elsevier, pp. 291–326.
- Messner, J.W., Mayr, G.J., Wilks, D.S. & Zeileis, A. (2014) Extending extended logistic regression: extended versus separate versus ordered versus censored. *Monthly Weather Review*, 142, 3003–3014.
- Miles, R. (1959) The complete amalgamation into blocks, by weighted means, of a finite set of real numbers. *Biometrika*, 46, 317–327.
- Molteni, F., Buizza, R., Palmer, T.N. & Petroliaigis, T. (1996) The ECMWF ensemble prediction system: methodology and validation. *Quarterly Journal of the Royal Meteorological Society*, 122, 73–119.
- Mösching, A. & Dümbgen, L. (2020a) Maximum likelihood estimation of a likelihood ratio ordered family of distributions. Preprint, arxiv.org/abs/2007.11521.
- Mösching, A. & Dümbgen, L. (2020b) Monotone least squares and isotonic quantiles. *Electronic Journal of Statistics*, 14, 24–49.
- Neelon, B. & Dunson, D.B. (2004) Bayesian isotonic regression and trend analysis. *Biometrics*, 60, 398–406.
- Newey, W.K. & Powell, J.L. (1987) Asymmetric least squares estimation and testing. *Econometrica*, 55, 819–847.
- Pappenberger, F., Ramos, M.H., Cloke, H.L., Wetterhall, F., Alfieri, L., Bogner, K. et al. (2015) How do I know if my forecasts are better? Using benchmarks in hydrologic ensemble prediction. *Journal of Hydrology*, 522, 697–713.
- R Core Team. (2020) *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Raftery, A.E., Gneiting, T., Balabdaoui, F. & Polakowski, M. (2005) Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review*, 133, 1155–1174.
- Rasp, S. & Lerch, S. (2018) Neural networks for postprocessing ensemble weather forecasts. *Monthly Weather Review*, 146, 3885–3900.
- Rigby, R.A. & Stasinopoulos, D.M. (2005) Generalized additive models for location, scale and shape (with discussion). *Journal of the Royal Statistical Society Series C*, 54, 507–554.
- Robertson, T. & Wright, F.T. (1975) Consistency in generalized isotonic regression. *Annals of Statistics*, 3, 350–362.
- Robertson, T. & Wright, F.T. (1980) Algorithms in order restricted statistical inference and the Cauchy mean value property. *Annals of Statistics*, 8, 645–651.
- Robertson, T., Wright, F.T. & Dykstra, R.L. (1988) *Order restricted statistical inference*. Hoboken: John Wiley & Sons.
- Rojo, J. & El Barmi, H. (2003) Estimation of distribution functions under second order stochastic dominance. *Statistica Sinica*, 13, 903–926.
- Rossi, B. (2013) Exchange rate predictability. *Journal of Economic Literature*, 51, 1063–1110.

- Schefzik, R., Thorarinsdottir, T.L. & Gneiting, T. (2013) Uncertainty quantification in complex simulation models using ensemble copula coupling. *Statistical Science* 28, 616–640.
- Scheuerer, M. (2014) Probabilistic quantitative precipitation forecasting using ensemble model output statistics. *Quarterly Journal of the Royal Meteorological Society*, 140, 1086–1096.
- Scheuerer, M. & Hamill, T.M. (2015) Statistical postprocessing of ensemble precipitation forecasts by fitting censored, shifted gamma distributions. *Monthly Weather Review*, 143, 4578–4596.
- Schlosser, L., Hothorn, T., Stauffer, R. & Zeileis, A. (2019) Distributional regression forests for probabilistic precipitation forecasting in complex terrain. *Annals of Applied Statistics*, 13, 1564–1589.
- Schulze Waltrup, L., Sobotka, F., Kneib, T. & Kauermann, G. (2015) Expectile and quantile regression — David and Goliath? *Statistical Modelling*, 15, 433–456.
- Seo, K. (2009) Ambiguity and second-order belief. *Econometrica*, 77, 1575–1605.
- Shaked, M. & Shanthikumar, J.G. (2007) *Stochastic orders*. New York: Springer.
- Shively, T.S., Sager, T.W. & Walker, S.G. (2009) A Bayesian approach to nonparametric monotone function estimation. *Journal of the Royal Statistical Society Series B*, 71, 159–175.
- Sloughter, J. M., Raftery, A.E., Gneiting, T. & Fraley, C. (2007) Probabilistic quantitative precipitation forecasting using Bayesian model averaging. *Monthly Weather Review*, 135, 3209–3220.
- Stellato, B., Banjac, G., Goulart, P. & Boyd, S. (2019) Osqp: quadratic programming solver using the ‘OSQP’ library. R package version 0.6.0.3.
- Stellato, B., Banjac, G., Goulart, P., Bemporad, A. & Boyd, S. (2020) OSQP: an operator splitting solver for quadratic programs. *Mathematical Programming Computation*, 12, 637–672.
- Stigler, S.M. (1975) The transition from point to distribution estimation. *Bulletin of the International Statistical Institute*, 46, 332–340.
- Swinbank, R., Kyouda, M., Buchanan, P., Froude, L., Hamill, T.M., Hewson, T.D. et al. (2016) The TIGGE project and its achievements. *Bulletin of the American Meteorological Society*, 97, 49–67.
- Taillardat, M., Mestre, O., Zamo, M. & Naveau, P. (2016) Calibrated ensemble forecasts using quantile regression forests and ensemble model output statistics. *Monthly Weather Review*, 144, 2375–2393.
- Taillardat, M., Fougères, A.L., Naveau, P. & Mestre, O. (2019) Forest-based and semiparametric methods for the postprocessing of rainfall ensemble forecasting. *Weather and Forecasting*, 34, 617–634.
- Tibshirani, J., Athey, S. & Wager, S. (2020) GRF: generalized random forests. R package version 1.2.0.
- Umlauf, N. & Kneib, T. (2018) A primer on Bayesian distributional regression. *Statistical Modelling*, 18, 219–247.
- Vannitsem, S., Wilks, D.S. & Messner, J. (Eds.) (2018) *Statistical postprocessing of ensemble forecasts*. Amsterdam: Elsevier.
- Vogel, P., Knippertz, P., Fink, A.H., Schlueter, A. & Gneiting, T. (2018) Skill of global raw and postprocessed ensemble predictions of rainfall over northern tropical Africa. *Weather and Forecasting*, 33, 369–388.
- Vovk, V., Shen, J., Manokhin, V. & Xie, M. (2019) Nonparametric predictive distributions based on conformal prediction. *Machine Learning*, 108, 445–474.
- Wilbur, W.J., Yeganova, L. & Kim, W. (2005) The synergy between PAV and AdaBoost. *Machine Learning*, 61, 71–103.
- Wilks, D.S. (2019) *Statistical methods in the atmospheric sciences*, 4th edn. Amsterdam: Elsevier.
- Winkler, R.L. (1972) A decision-theoretic approach to interval estimation. *Journal of the American Statistical Association*, 67, 187–191.
- Wu, J., Meyer, M.C. & Opsomer, J.D. (2015) Penalized isotonic regression. *Journal of Statistical Planning and Inference*, 161, 12–24.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section.

How to cite this article: Henzi A, Ziegel JF, Gneiting T. Isotonic distributional regression. *J R Stat Soc Series B*. 2021;83:963–993. <https://doi.org/10.1111/rssb.12450>