



Simplifying Random Forests' Probabilistic Forecasts

Nils Koster & Fabian Krüger

To cite this article: Nils Koster & Fabian Krüger (09 Oct 2025): Simplifying Random Forests' Probabilistic Forecasts, *The American Statistician*, DOI: [10.1080/00031305.2025.2552284](https://doi.org/10.1080/00031305.2025.2552284)

To link to this article: <https://doi.org/10.1080/00031305.2025.2552284>



© 2025 The Author(s). Published with
license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 09 Oct 2025.



[Submit your article to this journal](#)



Article views: 516



[View related articles](#)



[View Crossmark data](#)

CrossMark

Simplifying Random Forests' Probabilistic Forecasts

Nils Koster^{a,b} and Fabian Krüger^a

^aKarlsruhe Institute of Technology, Karlsruhe, Germany; ^bBroad Institute of MIT and Harvard, Cambridge, MA

ABSTRACT

Since their introduction by Breiman, Random Forests (RFs) have proven to be useful for both classification and regression tasks. The RF prediction of a previously unseen observation can be represented as a weighted sum of all training sample observations. This nearest-neighbor-type representation is useful, among other things, for constructing forecast distributions (as in Meinshausen's Quantile Regression Forests). In this article, we consider simplifying RF-based forecast distributions by sparsifying them. That is, we focus on a small subset of k nearest neighbors while setting the remaining weights to zero. This simplification, which we refer to as "Topk", greatly improves the interpretability of RF predictions. It can be applied to any forecasting task without re-training existing RF models. In empirical experiments, we document that the simplified predictions can be similar to or exceed the original ones in terms of forecasting performance. We explore the statistical sources of this finding via a stylized analytical model of RFs. The model suggests that simplification is particularly promising if the unknown true forecast distribution contains many small weights that are estimated imprecisely.

ARTICLE HISTORY

Received November 2024
Accepted August 2025

KEYWORDS

Forecast distribution;
Random forests; Simplicity

1. Introduction

Many statisticians agree that forecast distributions are preferable to mere point forecasts. Correspondingly, an active literature is concerned with making statistical forecast distributions in meteorology (e.g., Rasp and Lerch 2018), economics (e.g., Krüger, Clark, and Ravazzolo 2017), energy (e.g., Taieb, Taylor, and Hyndman 2021), epidemiology (e.g., Cramer et al. 2022), and other fields. Nevertheless, point predictions still dominate in many practical settings in policy, business, and society. As argued by Raftery (2016), the cognitive load that forecast distributions impose upon their users may be an important bottleneck impeding their adoption. Motivated by this possibility, we consider simplifying the forecast distributions produced by Random Forests (RFs; Breiman 2001; Meinshausen 2006), and study how simplification affects statistical forecasting performance. While our main focus is on probabilistic forecasting, the method we propose can also be used to simplify point forecasts for the mean.

More specifically, we approximate an RF forecast distribution for a continuous scalar outcome by a discrete distribution with k support points, where $k \in \{1, 2, \dots, n\}$ is a user-determined parameter and n denotes the number of training samples. The resulting forecast distribution can be cast as a collection of k scenarios, each occurring with a specified probability. If k is sufficiently small, this setup can effectively be communicated to non-statisticians. For example, Altig et al. (2022) use it to survey business executives about firm outcomes like future sales growth and employment. Abbas and Howard (2015, chap. 35) provide a textbook discussion from a decision analysis perspective.

Our approximation, which we refer to as "Topk" in the following, builds upon the fact that RFs can be cast as a nearest-neighbor-like method (Lin and Jeon 2006; Meinshausen 2006): The forecast distribution for a test sample observation with feature vector \mathbf{x}_0 is a discrete distribution with support points given by the training sample outcomes $(y_i)_{i=1}^n$ and corresponding probability weights $(w_i(\mathbf{x}_0))_{i=1}^n$. Meinshausen (2006) called these forecast distributions Quantile Regression Forests (QRFs). In the following, we do not distinguish between QRFs and RFs since the former can be constructed from the latter without any additional model fitting. As detailed in Section 2.1, the weight $w_i(\mathbf{x}_0)$ reflects the similarity between \mathbf{x}_0 and \mathbf{x}_i , the feature vector of the i th training sample observation. Consider a given test case \mathbf{x}_0 and the set $\mathcal{I}_k(\mathbf{x}_0)$ containing the indices of the k largest weights for this test case, where $k \in \{1, 2, \dots, n\}$. We then set

$$\tilde{w}_i(\mathbf{x}_0) = \begin{cases} \frac{w_i(\mathbf{x}_0)}{\sum_{j \in \mathcal{I}_k(\mathbf{x}_0)} w_j(\mathbf{x}_0)} & \text{if } i \in \mathcal{I}_k(\mathbf{x}_0) \\ 0 & \text{else} \end{cases}$$

That is, we retain only the $k << n$ largest weights, and rescale them such that they sum to one. All other weights are set to zero. Figure 1 provides a schematic description. Setting $k = n$ recovers the weights of the initial RF's mean and the QRF's distribution forecast, which is described in more detail in Section 2.¹ Larger values of k correspond to a larger number of scenarios. This increases the complexity of the forecast distribution, but may be beneficial in terms of statistical forecasting performance.

CONTACT Fabian Krüger  fabian.krueger@kit.edu  Karlsruhe Institute of Technology, Bluecherstrasse 17, 76185 Karlsruhe, Germany.

¹Prediction as in Equations (4) and (5) is unaltered, except for the substitution of $w_i(\mathbf{x}_0)$ by $\tilde{w}_i(\mathbf{x}_0)$.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/TAS.

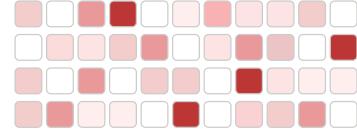
© 2025 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

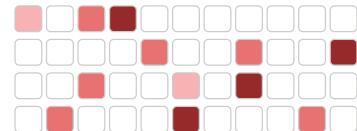
Step 1: Train Random Forest



Step 2: Compute Weight Matrix



Step 3: Select Top k=3 Weights & Re-Normalize



Step 4: Predict

$$y_1 + y_3 + y_4$$

⋮

$$y_2 + y_6 + y_{10}$$

Figure 1. Summary of proposed Top_k method. This schematic description summarizes the method proposed in this article. First, a standard RF is trained. Second, from the trained RF, we compute a weight vector for each new test case. Collecting the weight vectors for multiple test cases results in a matrix, with rows corresponding to test cases and columns corresponding to training cases. Next, we select the top k (in this case $k = 3$) weights for each test case, re-normalize such that the weights again sum up to one and set the remaining weights to zero. These weights can now be used for prediction, as illustrated here for the first and fourth test cases.

Practical Illustration. We provide a practical illustration of our method by considering labor market survey data provided by DIW (2022). The dataset is an openly available subsample of the German Socioeconomic Panel (SOEP; Goebel et al. 2019). It is recorded between 2015 and 2019, covering 21,851 observations corresponding to 5847 distinct persons after preprocessing. Many participants were asked repeatedly over the survey years, so that the data has a panel structure. Our goal here is to predict a person’s annual salary in Euros based on socioeconomic regressors such as age, education, family status, and the industry in which the person is employed. We use data from 2015 to 2018 for training and 2019 for testing, and consider a Top_k model with $k = 5$ along with a full RF. Compared to the full RF, the Top_k model is about 6% worse in terms of distribution forecasting performance (as measured by the CRPS, introduced in Section 2.2), and about 3% better in terms of point forecasting performance (as measured by the squared error). Section A of the online supplement provides details on the dataset and RF forecasting performance.

To illustrate, consider a fully employed 47-year-old male, living alone and working in the financial sector. We aim to predict his salary in 2019. A full RF predicts a mean income of 84.705, placing nonzero weight on 171 training sample observations. The simplified Top_5 model uses only the five largest of these 171 weights which (before re-normalization) sum up to 0.47. To re-normalize, we simply multiply the five remaining weights by $1/0.47$. The five support points (the training sample observations that correspond to the remaining weights) are listed in Table 1, together with their weights. With those, the Top_5 model predicts a mean income of 84.184.

For communicating these results, the five remaining weights and support points can be cast as scenarios, spanning an income range from 51.608 to 100.433 Euros. The most likely scenario,

Table 1. SOEP test prediction.

	Survey Year	Female	Age	No. Per.	No. Child.	Years Educ.	Empl.	Sector ID	$w_j(x_0)$ [%]	Income
x_0	2019	False	47	1	0	14.5	F	64	–	83.279
x_{51}	2017	False	45	1	0	14.5	F	64	30.6	94.903
x_{52}	2018	False	45	1	0	18	F	64	24.6	79.206
x_{53}	2017	False	44	1	0	15	F	64	17.3	100.433
x_{54}	2016	False	44	1	0	14.5	F	64	16.0	51.608
x_{55}	2015	False	43	1	0	14.5	F	64	11.6	87.173

NOTE: The table shows the five support points x_{51} to x_{55} for a test point (first row, x_0) in the Top_5 model. “Empl.” abbreviates the variable “employed”, “F” abbreviates “full employment”. Sector ID 64 corresponds to “Provision of financial services”. The corresponding (rounded) weight of each support point is depicted in the penultimate column, while we report the respective income in the last column.

occurring with a probability of 30.6%, involves an income of 94.903 Euros. Due to the structure of RF forecast distributions, the selected scenarios are directly associated with five actual observations occurring in our training sample. In particular, three of the five scenarios (number 1, 4, and 5) listed in Table 1 involve responses of the same individual in previous years. This situation is natural since some of the regressors are nearly constant over time, so that the test point x_0 is likely to be similar to training data x_i corresponding to the same individual. Using the same individual’s past incomes also seems plausible from a practical perspective.² Table 1 also yields relevant (albeit implicit) substantive information: For example, the fact that all five scenarios involve individuals of the same gender, family status and industry sector indicate that the RF model considers these regressors highly relevant for prediction. Conversely, the five scenarios cover all survey years present in the training data

²If one wanted to avoid the use of data from the same individual, one could reduce the dataset to a single observation per individual, and otherwise use the same methodology.

(2015–2018), indicating that time variation in incomes is less important.

Related Literature. We approach the interpretability of a RF from a different angle compared to most other existing literature on the topic: In Breiman’s original work on RFs, feature permutation was introduced, effectively quantifying the influence of each feature on the overall performance of the model. This method is still widely used and implemented in standard software packages (Pedregosa et al. 2011). Biau and Scornet (2016, sec. 5) provide further discussion. This perspective on interpretation is also commonly used for other forecasting models, such as neural networks, see for example Lundberg and Lee (2017). Visualizations of RFs have also been proposed (Zhao et al. 2019; Haddouchi and Berrado 2019). In addition, alternative notions of sparsity have been considered in the context of tree-type models. For example, Nan, Wang, and Saligrama (2016) use the notion of sparsity in terms of feature usage in order to speed up predictions. Several other studies consider combinations of simple prediction rules, building upon ideas in Friedman and Popescu (2008). The “node harvest” method by Meinshausen (2010) is particularly interesting since it can be cast as a weighted empirical distribution, and can thus be used for probabilistic prediction. Node harvest first generates a large number of “nodes”, each of which defines a simple subspace of the predictor space. It then uses a constrained optimization problem to select the nodes that are most useful for prediction. Although sparsity is not enforced, Meinshausen finds that the optimal solution often involves a small number of nodes only. However, this type of sparsity (with respect to nodes) is different from the type of sparsity we consider, which is with respect to the number of training sample observations being considered for prediction. Indeed, these two notions of sparsity seem highly complementary: A forecast distribution constructed via node harvest may involve a large number of training sample observations, that is be far from sparse in our sense. Conversely, a forecast distribution constructed using our proposed methodology implicitly involves many regression trees, thus being far from sparse in Meinshausen’s sense.

Finally, many extensions and modifications of Breiman’s original RF have been proposed in the literature, typically with the aim of improving statistical performance (either empirically, or in terms of theoretical properties). See Biau and Scornet (2016) for a survey, and for example Beck, Kozbur, and Wolf (2023) for a forecast combination perspective, Cevid et al. (2022) for a multivariate model that links RFs to kernel-based methods, and Wager and Athey (2018) for an adaption to causal inference. By contrast, we focus on a standard, univariate RF implementation, and post-process its forecast distribution in a way that makes it easier to interpret. The procedure we propose can easily be applied to other RF variants (or even to prediction methods other than RFs), provided that they can be represented as discrete distributions of the type described above.

Roadmap of the Article. The remainder of this article is organized as follows: Section 2 introduces our methodological setup, including RFs and methods for evaluating forecast distri-

butions. In Section 3, we study the performance of our ‘Top k ’ simplification in a series of empirical experiments based on 18 datasets for which RFs have been found to perform well, as compared to deep neural network models (Grinsztajn, Oyallon, and Varoquaux 2022). Our findings indicate that already for $k = \{5, 10\}$, the simplified RFs can perform on a similar level compared to the full counterpart for both probabilistic and point forecasts, depending on the dataset. Considering probabilistic forecasts, for $k = 20$, the median performance across all datasets is equivalent to the full RFs and considering $k = 50$ even increases the median performance slightly. For point (mean) forecasts, $k = 20$ even increases performance slightly. Even though very sparse choices like $k \in \{3, 5\}$ may come with a significant performance decrease, we argue that they may be worthwhile if ease of communication is a main concern. We further show that our results are qualitatively robust to different hyperparameter choices. In order to rationalize the empirical results, Section 4 then considers a detailed analytical example that models the weights estimated by RFs as a random draw from a Dirichlet-type distribution. The example also features a true vector of weights that may be either similar to, or different from, the estimated weights. This setup is useful to study how a simplifying approximation similar to Top k affects statistical forecasting performance. In a nutshell, the amount of noise in the estimated weights determines whether or not the simplification comes at a high cost in terms of performance. Perhaps surprisingly, one can construct examples in which simplification improves performance: this result arises when the largest weights are estimated precisely, whereas smaller weights are more noisy. Focusing on the largest weights then constitutes a beneficial form of shrinkage. When the small weights are estimated precisely, however, simplification is harmful in terms of performance. The online supplement contains further empirical results and detailed derivations for Section 4. Replication materials are available at https://github.com/kosnil/simplify_rf_dist.

2. Methodological Setup

In this section, we describe RF based forecasting as well as the forecast evaluation methods we consider.

2.1. Forecasting Methods

Here we describe Random Forests and their probabilistic cousins, Quantile Regression Forests. We follow Lin and Jeon (2006) and Meinshausen (2006) who emphasize the perspective of RF predictions as a weighted sum over training observations. We refer to Hastie, Tibshirani, and Friedman (2009) for a textbook presentation of RFs.

Our goal is to fit a univariate forecasting model. That is, we have some dataset $\mathcal{D} = (\mathbf{x}_i, y_i)_{i=1}^n$ of training set size n , where \mathbf{x}_i is a p -dimensional vector of features, and $y_i \in \mathbb{R}$ is a real-valued outcome. We use this dataset to train our model \hat{f} , such that some choice of loss function \mathcal{L}_n is minimized:

$$\min_{\hat{f}} \mathcal{L}_n(\hat{f}) = \frac{1}{n} \sum_i^n \mathcal{L}(\hat{f}(\mathbf{x}_i), y_i).$$

In a typical regression context, the function $\hat{f} : \mathbb{R}^p \rightarrow \mathbb{R}$ maps \mathbf{x}_i to the real line in order to estimate the conditional expectation functional. A popular loss function is given by squared error, with $\mathcal{L}(z, y) = (y - z)^2$. Let further $\mathcal{B} \subseteq \mathbb{R}^p$ denote the feature space, that is, the space in which the individual input samples \mathbf{x}_i exist.

Random Forests. An RF is an ensemble of individual regression trees, each denoted by $\mathcal{T}(\xi)$, where ξ describes the configuration of the tree. At each node, a single tree greedily splits \mathcal{B} and rectangular subspaces thereof into two further rectangular subspaces, such that the loss \mathcal{L}_n is minimized. Each resulting subspace corresponds to a leaf $\mathcal{R}_l \subseteq \mathcal{B}$, $l = 1, \dots, L$ where L is the total number of leaves. Each sample \mathbf{x}_i can only occur in one leaf or, put differently, when dropping a sample \mathbf{x}_i down the tree, it can only fall into one leaf. This leaf is denoted $\ell(\mathbf{x}_i, \xi) \in \{1, \dots, L\}$ for tree $\mathcal{T}(\xi)$. For a single tree $\mathcal{T}(\xi)$, a prediction $\hat{\mu}_{\mathcal{T}}(\mathbf{x}_0)$ for a new sample \mathbf{x}_0 is obtained by taking the mean of all training samples within leaf $\ell(\mathbf{x}_0, \xi)$. This can be expressed as

$$\hat{\mu}_{\mathcal{T}}(\mathbf{x}_0) = \sum_{i=1}^n w_i(\mathbf{x}_0, \xi) y_i, \quad (1)$$

where the weight $w_i(\mathbf{x}_0, \xi)$ is equal to zero for all training samples i that fall into leaves other than $\ell(\mathbf{x}_0, \xi)$, and is equal to one over the leaf size for all training samples that fall into $\ell(\mathbf{x}_0, \xi)$:

$$w_i(\mathbf{x}_0, \xi) = \frac{\mathbb{1}\{\mathbf{x}_i \in \mathcal{R}_{\ell(\mathbf{x}_0, \xi)}\}}{\sum_{j=1}^n \mathbb{1}\{\mathbf{x}_j \in \mathcal{R}_{\ell(\mathbf{x}_0, \xi)}\}}. \quad (2)$$

Motivated by the lack of stability and tendency to overfit of individual trees, RFs build B trees $(\mathcal{T}(\xi_b))_{b=1}^B$, based on B bootstrap samples of \mathcal{D} , and consider their average prediction. Moreover, in each split within each tree, it is common to consider only a random subsample of \tilde{p} out of p regressors. This step aims to diversify the ensemble of trees by avoiding excessive use of the same regressors for splitting. Common choices for \tilde{p} are $\lfloor \sqrt{p} \rfloor$ or $\lfloor \frac{p}{3} \rfloor$ (Probst, Wright, and Boulesteix 2019), where $\lfloor z \rfloor$ floors the real number z to the nearest integer. The RF mean prediction can thus be expressed as

$$\begin{aligned} \hat{\mu}_{\text{RF}}(\mathbf{x}_0) &= \frac{1}{B} \sum_{b=1}^B \sum_{i=1}^n w_i(\mathbf{x}_0, \xi_b) y_i \\ &= \sum_{i=1}^n w_i(\mathbf{x}_0) y_i, \end{aligned} \quad (3)$$

where

$$w_i(\mathbf{x}_0) = \frac{1}{B} \sum_{b=1}^B w_i(\mathbf{x}_0, \xi_b) \quad (4)$$

is the weight for training sample i , averaged across all B trees. By construction, the weights $(w_i(\mathbf{x}_0))_{i=1}^n$ are nonnegative and sum to one. Thus, $w_i(\mathbf{x}_0)$ can be interpreted as the empirically estimated probability that the new test sample observation is equal to y_i .

Quantile Regression Forest. Conceptually, $\hat{\mu}_{\text{RF}}(\mathbf{x}_0)$ is an estimate of the conditional mean $\mathbb{E}[Y|X = \mathbf{x}_0]$. As described above, it is obtained as a weighted sum over all training observations. Meinshausen (2006) extends this framework to estimating the cumulative distribution function (CDF) of Y , which is given by $\mathbb{E}[\mathbb{1}(Y \leq t)|X = \mathbf{x}_0] = \mathbb{P}(Y \leq t|X = \mathbf{x}_0)$, where $t \in \mathbb{R}$ is a threshold value. The similarity to RFs becomes apparent in the last expression. Using the weights from (4), one can approximate the CDF by the weighted mean over the binary observations $\mathbb{1}(y_i \leq t)$:

$$\hat{\mathbb{P}}(Y \leq t|X = \mathbf{x}_0) = \sum_{i=1}^n w_i(\mathbf{x}_0) \mathbb{1}(y_i \leq t). \quad (5)$$

That is, QRFs estimate the CDF of Y via the weighted empirical CDF of the training sample outcomes $(y_i)_{i=1}^n$, using the weights $(w_i(\mathbf{x}_0))_{i=1}^n$ produced by RFs. This estimator is practically appealing as it arises as a byproduct of standard RF software implementation. Furthermore, its representation in terms of a weighted empirical CDF enables a theoretical understanding of its properties by leveraging tools from nonparametric statistics (Lin and Jeon 2006; Meinshausen 2006). In this article, we consider the standard variant of QRFs which uses squared error as a criterion for finding splits (and thus growing the forest's individual trees). Various other splitting criteria have been analyzed in the literature. In particular, Cevid et al. (2022) propose to use a splitting criterion based on distributional similarity. Since their RF variant retains the weighted empirical CDF representation (see their Section 2.2), our Topk method can be applied to it as well.

2.2. Forecast Evaluation

Since we generate probabilistic forecasts, we need a tool to evaluate them. For this, we use the Continuous Ranked Probability Score (CRPS), a strictly proper scoring rule. Scoring rules are loss functions for probabilistic forecasts. We use them in negative orientation, so that smaller scores indicate better forecasts. When evaluated using a proper scoring rule, a forecaster minimizes their expected score by stating what they think is the true forecast distribution. Under a strictly proper scoring rule, this minimum is unique within a suitable class of forecast distributions. Conceptually, strictly proper scoring rules incentivize careful and honest forecasting. See Gneiting and Raftery (2007) for a comprehensive technical treatment, and Winkler (1996) and Gneiting and Katzfuss (2014) for further discussion and illustration. The CRPS (Matheson and Winkler 1976) is defined as

$$\text{CRPS}(\hat{F}, y) = \int_{-\infty}^{\infty} \left(\hat{F}(z) - \mathbb{1}\{z \geq y\} \right)^2 dz \quad (6)$$

where \hat{F} denotes the CDF implied by the forecast distribution and y denotes the true outcome. Gneiting and Raftery (2007, sec. 4.2) discuss connections to the concept of statistical energy (e.g., Székely and Rizzo 2017) and note that the CRPS can equivalently be expressed via its kernel representation given by $\text{CRPS}(\hat{F}, y) = \mathbb{E}_{\hat{F}}[|Z - y|] - \frac{1}{2} \mathbb{E}_{\hat{F}}[|Z - Z'|]$, where Z and Z' are two independent draws from \hat{F} . For various forms of

\hat{F} , efficient computation methods for the CRPS are available, either via analytical expressions (if \hat{F} is parametric) or via a representation due to Laio and Tamea (2007) (if \hat{F} is a discrete empirical distribution). See Jordan, Krüger, and Lerch (2019) for an overview, and Jordan (2016) for details.

The CRPS allows for very general types of forecast distributions \hat{F} . In particular, the forecast distribution may be discrete, that is, it need not possess a density. This allows for evaluating forecast distributions based on (weighted) empirical CDFs, which arise in the case of QRFs. In the special case that the forecast distribution is deterministic, that is, it places point mass on a single outcome, the CRPS reduces to the Absolute Error (AE). Thus, numerical values of the AE and CRPS can meaningfully be compared to each other.

Jordan, Krüger, and Lerch (2019) provide an efficient implementation of the CRPS for weighted empirical distributions in their R-package `scoringRules`. Let $(y_i)_{i=1}^n$ denote the response values from the training data, and denote by $y_{(i)}$ their i th ordered value, with $y_{(1)} \leq y_{(2)} \leq \dots \leq y_{(n)}$. Furthermore, let $w_{(i)}$ denote the weight corresponding to $y_{(i)}$. Then the CRPS for a realization $y \in \mathbb{R}$ is given by

$$\text{CRPS}(\hat{F}, y) = 2 \sum_{i=1}^n w_{(i)} (y_{(i)} - y) \times \left(\mathbb{1}\{y < y_{(i)}\} - \left(\sum_{j=1}^i w_{(j)} \right) + \frac{w_{(i)}}{2} \right), \quad (7)$$

where we dropped the dependence of $w_{(i)}$ on a vector \mathbf{x}_0 of covariates at this point for ease of notation. Equation (7) extends Jordan, Krüger, and Lerch's eq. (3) to the case of non-equal weights, based on their implementation in the function `crps_sample`. In the case of sparse weights, one may omit the indices i with $w_{(i)} = 0$ from the sum at (7) in order to speed up the computation.

Additionally to the CRPS, we also report results for the squared error (SE). In the present context, the SE is given by

$$\text{SE}(\hat{F}, y) = (y - \sum_{i=1}^n w_i y_i)^2. \quad (8)$$

Hence, the SE depends on the forecast distribution \hat{F} via its mean $\sum_{i=1}^n w_i y_i$ only.

3. Experimental Results

In order to assess the statistical performance of the simplified forecast distributions, we conduct experiments on 18 datasets considered by Grinsztajn, Oyallon, and Varoquaux (2022) in the context of numerical regression. The authors demonstrate that tree-based methods compare favorably to neural networks for these datasets. Their selection of datasets aims to represent real-world, "clean" datasets with medium size as well as heterogeneous data types and fields of applications. If deemed necessary, some basic preprocessing was applied by Grinsztajn, Oyallon, and Varoquaux (2022). Details can be found in Section 3.5 and Appendix A.1 in their paper, and in Table S3 of our online supplement. The dataset `delays_zurich_transport` contains about 5.6 million data points in its original form. For computational reasons, we reduced the size of this dataset through

Table 2. Results for forecast distributions.

Dataset	Absolute CRPS Full	CRPS relative to full				
		Top3	Top5	Top10	Top20	Top50
cpu_act	1.2508	1.21	1.08	1.00	0.96	0.95
pol	1.4203	1.50	1.28	1.10	1.01	0.96
elevators	0.0014	1.22	1.10	1.01	0.96	0.95
wine_quality	0.2565	1.35	1.20	1.09	1.02	0.99
Ailerons	0.0001	1.23	1.11	1.01	0.97	0.96
houses	0.1229	1.19	1.08	0.99	0.96	0.95
house_16H	0.1984	1.32	1.17	1.06	1.01	0.98
diamonds	0.1320	1.34	1.21	1.11	1.05	1.01
Brazilian_houses	0.0256	0.88	0.82	0.78	0.79	0.85
Bike_Sharing_Demand	46.0292	1.26	1.15	1.06	1.02	1.00
nyc-taxi-green-dec-2016	0.1505	1.39	1.24	1.12	1.06	1.02
house_sales	0.0995	1.25	1.14	1.04	0.99	0.97
sulfur	0.0123	1.05	0.97	0.94	0.95	0.97
medical_charges	0.0376	1.35	1.21	1.11	1.05	1.01
MiamiHousing2016	0.0778	1.20	1.10	1.02	0.98	0.97
superconduct	3.6341	1.14	1.08	1.02	1.00	0.99
yprop_4_1	0.0140	1.38	1.25	1.14	1.08	1.03
delays_zurich_transport	1.6174	1.33	1.19	1.10	1.05	1.02
Median	—	1.25	1.14	1.05	1.00	0.98

NOTE: The table reports the CRPS of the full RF as well as the CRPS for $\text{Top}\{3, 5, 10, 20, 50\}$ relative to the full RF, that is, $\text{CRPS}_{\text{Top } k} / \text{CRPS}_{\text{Full}}$. A value smaller than 1 means that $\text{Top } k$ outperforms the full RF. The last row lists the median relative CRPS across datasets.

random subsampling, to approximately 1.1 million data points (20% of the original observations). We did not apply further preprocessing of any of the datasets in order to retain comparability. We allocate 70% of each dataset for training our models and reserve the remaining 30% for testing. For each dataset, we train a RF, and then evaluate its performance by computing the average CRPS and SE, as introduced in (7) and (8) over the test dataset. Our main interest is in studying the impact of the parameter k which governs the number of support points of the sparsified forecast distribution. We consider a grid of choices $k = 1, \dots, 50$ and denote these sparse RFs as "Top k ". Our standard choice of RF hyperparameters is a combination of default values of the machine learning software packages `scikit-learn` (Pedregosa et al. 2011) and `ranger` (Wright and Ziegler 2017) which are popular choices in the Python (van Rossum et al. 2011) and R (R Core Team 2022) programming languages, respectively. In summary, we consider a random selection of \sqrt{p} out of p possible features at each split point, do not impose any form of regularization in terms of tree growth restriction, and set the number of trees to 1000 in order to obtain a large and stable ensemble. These choices, which are also listed in the first row (entitled "standard") of Table S5 in the online supplement, are used for the analysis in Sections 3.1 and 3.2. In Section 3.3, we further consider the effects of tuning hyperparameters.

3.1. Probabilistic Forecasts

Table 2 presents the CRPS for the full RF and the relative CRPS of $\text{Top } k$, for $k \in \{3, 5, 10, 20, 50\}$, compared to the full RF. For example, a relative CRPS of 1.5 indicates a 50% larger CRPS for the $\text{Top } k$ version compared to the full RF. The three smallest values for k seem especially attractive in terms of simplicity and ease of communication. While the two larger choices $k \in \{20, 50\}$ are less attractive in terms of simplicity, we also consider them in order to assess the tradeoff between simplicity and performance.

Using only $k = 3$ support points performs worse than the full RF, with a median performance cost of 25% across datasets. While this result is unsurprising from a qualitative perspective, its magnitude is interesting, and gives a first indication of the performance cost of using a rather drastic simplification of the original forecast distribution. For each single dataset, we find that the performance of Top_k improves monotonically when increasing k from 3 to 5, from 5 to 10, and for all datasets but one when increasing k from 10 to 20. This pattern is plausible, given that we move from a drastic simplification ($k = 3$) to less drastic versions. Compared to the full RF, Top_5 implies a median loss increase of 14%, whereas Top_{10} yields a median loss increase of 5%. Top_{20} performs equally well as the full RF in the median. For most datasets, Top_{50} slightly enhances predictive performance compared to the full RF. Whether $k = 50$ support points remain worth interpreting depends on the application at hand. Only for a few datasets, Top_{50} is not sufficient to reach the performance of the full RF, but performance costs are small even for these datasets.

In order to contextualize the magnitude of our presented results (such as Top_3 's median CRPS increase of 25% compared to the full RF), we next present results on two simple benchmark methods.³ First, we consider a deterministic point forecast which assumes that the full RF's median forecast materializes with probability one.⁴

This is a very optimistic point (rather than probabilistic) forecast, containing no uncertainty. As noted earlier, its CRPS is the same as its AE. We consider this benchmark in order to quantify the costs of ignoring uncertainty altogether. We clearly expect these costs to be positive, that is, we expect the point forecast to perform worse than the full RF. Second, we use the CRPS of the unconditional empirical distribution of the response variable in the training sample. This distribution places a uniform weight of $1/n$ on each training sample observation, in contrast to the QRF weights $w(\mathbf{x}_0)$ that depend on the feature vector \mathbf{x}_0 . The unconditional distribution is a very conservative forecast, as no information about the features is used whatsoever. Qualitatively, we clearly expect this forecast to perform worse than the full RF. Quantitatively, the difference in performance of the unconditional versus conditional forecast distributions captures the predictive content of the features (see e.g., Gneiting and Resin 2023, sec. 2.5). Both benchmarks are visualized jointly with Top_3 results in Figure 2.

As expected, the relative CRPS of the point forecast (shown in blue) exceeds one for most datasets, indicating that it is generally inferior to the full RF model. Notably, an exception is observed for `Brazilian_houses`, where the point forecast's relative CRPS is smaller than one. This result appears to be due to prediction uncertainty being very small for this dataset; see Figure 13 in the online supplement for Grinsztajn, Oyallon, and

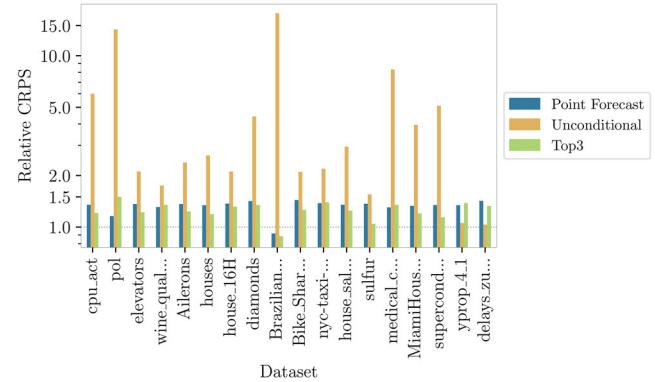


Figure 2. Benchmarking performance. Blue bars: Relative CRPS of a deterministic point forecast (median) as a benchmark that ignores uncertainty. Ochre bars: Relative CRPS of the unconditional forecast distribution, yielding a conservative benchmark that ignores features. Green bars: Relative CRPS of Top_3 , as listed in Table 2. In each case, relative CRPS results are compared to the full RF. That is, a relative CRPS smaller than 1 (represented by horizontal line) indicates that the method performs better than the full RF.

Varoquaux (2022). Indeed, for this dataset, the response variable seems to mostly be a linear combination of a subset of the features. Furthermore, Top_3 (shown in green) performs similar to or better than the point forecast for all datasets, demonstrating the usefulness of incorporating additional uncertainty in the forecast. Note that the point forecast has access to the full RF forecast distribution, using the median of this distribution as a point forecast. By contrast, Top_3 only has access to the three most important support points of the RF forecast distribution.

The relative CRPS of the unconditional forecast distribution (displayed in ochre in Figure 2) exceeds two for most datasets, indicating that the features are generally very useful for prediction. An exception occurs for the last two datasets (`yprop_4_1` and `delays_zurich_transport`). In these cases, the RF appears unable to learn meaningful connections between the features and the target variable. Figure 13 in the online supplement of Grinsztajn, Oyallon, and Varoquaux (2022) supports this interpretation, reporting low predictability (in terms of low out-of-sample R^2) for these datasets. In this situation, we cannot expect a Top_k -model to perform well compared to the full RF. To see this, consider the stylized case of the features being entirely uninformative. Subsequently, the unconditional distribution (placing a weight of $1/n$ on all training sample responses) is the best possible forecast, which is in sharp contrast to Top_k (for which k weights are nonzero and large by construction, whereas the remaining $n - k$ weights are forced to zero).

In order to further study the properties of Top_k forecast distributions, Table 3 reports the average weight sums across test cases, for different values of k . The weight sums are computed before our normalization step (see Section 1) which re-scales all weight sums to one. For a given choice of k , the unnormalized weight sums can vary in magnitude, both across datasets and from test case to test case. By construction, the weight sums increase with k . Interestingly, many datasets yield a large weight sum for Top_3 , exceeding 10% for all but four datasets. For Top_{50} , the average weight sum exceeds 50% for most datasets. These numbers are remarkable, given that the datasets include thousands of training samples (n , see rightmost column of Table 3) that could potentially be used as support points for the RF

³We refer to Grinsztajn, Oyallon, and Varoquaux (2022) for detailed comparisons of RF point forecasts to other tree-based models and neural networks.

⁴Formally, this forecast is characterized by the CDF $\hat{F}_\delta(z) = \mathbb{1}(z \geq \text{med}(\hat{F}_{\text{full}}(z)))$, where $\text{med}(\hat{F}_{\text{full}})$ denotes the median implied by the CDF of the full RF's forecast distribution. That is, the CDF \hat{F}_δ is a step function with a single jump point at the median forecast of the full RF. We choose the median functional here because the latter is the optimal point forecast under absolute error loss, to which the CRPS reduces in the case of a deterministic forecast.

Table 3. Top k weight sums.

Dataset	Average sum					n
	Top3	Top5	Top10	Top20	Top50	
cpu_act	0.094	0.135	0.212	0.318	0.504	5734
pol	0.073	0.104	0.159	0.233	0.359	10.500
elevators	0.119	0.166	0.252	0.366	0.555	11.619
wine_quality	0.150	0.189	0.258	0.350	0.513	4547
Ailerons	0.119	0.168	0.257	0.374	0.566	9625
houses	0.143	0.201	0.304	0.435	0.634	14.447
house_16H	0.071	0.102	0.162	0.247	0.402	15.948
diamonds	0.257	0.348	0.494	0.656	0.851	37.758
Brazilian_houses	0.202	0.278	0.406	0.558	0.763	7484
Bike_Sharing_Demand	0.248	0.334	0.473	0.628	0.821	12.165
nyc-taxi-green-dec-2016	0.172	0.234	0.337	0.461	0.642	407.284
house_sales	0.116	0.160	0.240	0.345	0.522	15.129
sulfur	0.212	0.296	0.438	0.602	0.811	7056
medical_charges	0.223	0.304	0.438	0.590	0.789	114.145
MiamiHousing2016	0.196	0.267	0.385	0.520	0.704	9752
superconduct	0.390	0.497	0.617	0.708	0.805	14.884
yprop_4_1	0.111	0.149	0.216	0.304	0.456	6219
delays_zurich_transport	0.015	0.024	0.048	0.095	0.230	765.180

NOTE: Average sum of un-normalized Top{3,5,10,20,50} weights. The last column reports the number of observations in the training set.

Table 4. Results for conditional mean forecasts.

Dataset	Absolute SE Full	SE relative to Full				
		Top3	Top5	Top10	Top20	Top50
cpu_act	6.5359	1.11	0.99	0.93	0.91	0.91
pol	38.5033	1.61	1.37	1.16	1.03	0.95
elevators	9.22×10^{-6}	1.09	0.98	0.89	0.85	0.86
wine_quality	0.3485	1.42	1.24	1.11	1.03	1.00
Ailerons	3.22×10^{-8}	1.15	1.02	0.94	0.91	0.91
houses	0.0590	1.16	1.03	0.94	0.91	0.92
house_16H	0.3011	1.49	1.31	1.15	1.06	1.00
diamonds	0.0563	1.37	1.24	1.12	1.06	1.02
Brazilian_houses	0.0072	1.05	1.12	0.99	0.94	0.94
Bike_Sharing_Demand	10126.7983	1.22	1.11	1.03	0.99	0.99
nyc-taxi-green-dec-2016	0.1528	1.38	1.23	1.12	1.05	1.02
house_sales	0.0382	1.22	1.11	1.02	0.96	0.94
sulfur	0.0015	0.69	0.69	0.73	0.81	0.90
medical_charges	0.0073	1.36	1.22	1.11	1.05	1.01
MiamiHousing2016	0.0243	1.15	1.05	0.97	0.94	0.94
superconduct	87.2592	1.05	1.02	0.97	0.95	0.95
yprop_4_1	0.0010	1.26	1.16	1.10	1.05	1.02
delays_zurich_transport	9.3282	1.33	1.19	1.10	1.05	1.02
Median	—	1.22	1.12	1.02	0.98	0.95

NOTE: The table reports the SE of the full RF as well as the SE for Top{3,5,10,20,50} relative to the full RF, that is, SE_{Topk}/SE_{Full} . The last row lists the median relative SE for each choice of k across datasets.

forecast distributions. If the weights were uniform, we would hence observe Top k weight sums of k/n . This is in sharp contrast to our empirical finding that a few large weights dominate for most datasets. Lin and Jeon (2006) find similar results in their work, where they consider RFs as adaptive nearest-neighbor methods and investigate the influence of the minimum number of samples per leaf. Figures 1(c), (d), and 5 show few large weights for synthetic datasets. The presence of a small number of important weights explains why the simplification pursued by Top k often results in modest (if any) performance costs as compared to the full RF.

3.2. Mean Forecasts

Let us turn our attention towards conditional mean forecasts, for which results in terms of squared error are shown in Table 4. We notice a similar pattern as in the probabilistic scenario: apart

from Top3 outperforming the full RF for one dataset (sulfur), Top3 performs worse than the full RF for the remaining datasets, with a maximum performance cost of 61% for pol.

This results in a median SE increase of 22% for Top3. Top5 still shows a 12% median increase and for Top10, the median performance almost matches the full RF's performance. Top20 and Top50 even outperform the full RF, yielding median improvements of 2% and 5%, respectively. Compared to the results for forecast distributions (Table 2), the results in Table 4 indicate that the performance costs of simplicity are comparatively lower in the case of mean forecasts, with median relative losses being somewhat smaller for a given value of k .

3.3. Varying Hyperparameters

Compared to other modeling algorithms, RFs have relatively few hyperparameters. Nevertheless, tuning its hyperparameters can improve the performance of RFs (Probst, Wright, and Boulesteix 2019). The most important hyperparameters control the depth of each individual tree, as well as the number of randomly selected regressors considered for splitting. The depth of a tree can be restricted directly ("maximum depth") or indirectly by restricting leaf and split sizes ("minimum leaf size" and "minimum split size," respectively). The number of regressors considered for splitting is often denoted as "max features" or "mtry" (the latter term is used, for example in the R packages `randomForest` (Liaw and Wiener 2002) and `ranger` (Wright and Ziegler 2017)). In what follows, we study how different hyperparameter sets influence the Top k prediction and whether a hyperparameter set that optimizes the full RF is also beneficial for Top k . We therefore investigate the influence of "max features" and one of the depth-regularizing hyperparameters, "minimum leaf size", on the Top k approach. To do so, we consider a grid search for both, the full RF and Top3, with 5-fold cross-validation on the training set of each dataset. Due to the size of `medical_charges`, `nyc-taxi-green-dec-2016` and `delays_zurich_transport`, we use a validation set which contains 25% of the training set instead. Further, the latter two are down-sampled to 30% and 15% of their original training set size, respectively. For brevity, our analysis of hyperparameter tuning focusses on the case $k = 3$, which is the most drastic simplification we consider.

Figure 3 visualizes the CRPS of Top3 relative to the full RF's CRPS for three different hyperparameter sets. In each case, we consider the same hyperparameter set for Top3 and the full RF. The green bars show the results with standard hyperparameters, as listed in Table 2. The red bars indicate the relative performance with the respective hyperparameter set that optimizes the full RF, while the blue bars visualize the relative performance with the hyperparameters that optimize the CRPS of Top3. When hyperparameters are tuned on the full RF, Top3 performance tends to slightly decrease overall. Across the datasets, the median relative CRPS is 1.31, compared to 1.25 for the standard setting. Conversely, if hyperparameters are tuned on Top3, the relative performance of Top3 is mostly better than in the standard version, reaching a median relative CRPS of 1.20.

Figure 4 presents results for point forecasting performance, which are similar to the probabilistic case. Tuning on the full RF

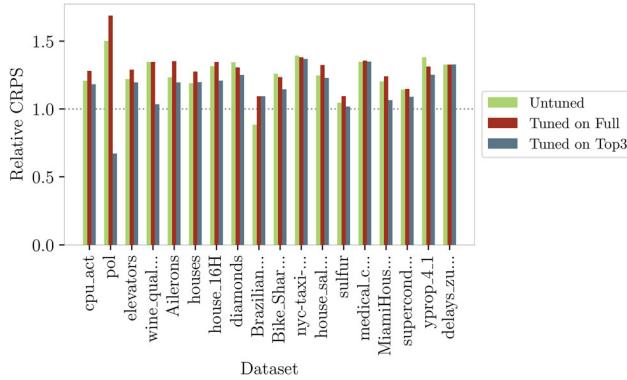


Figure 3. Hyperparameter tuning performance (probabilistic forecasts). The figure shows the relative CRPS of Top3 compared to the full RF, for different hyperparameter settings. Green bars: standard hyperparameters, as listed in Table 2. Red bars: hyperparameters that optimize the full RF. Blue bars: hyperparameters that optimize Top3. Hyperparameter tuning is based on CRPS.

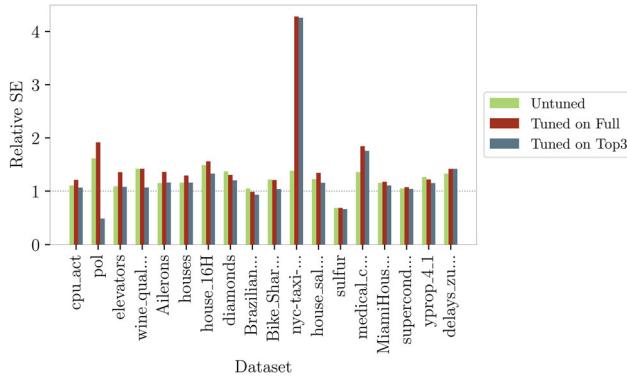


Figure 4. Hyperparameter tuning performance (point forecasts). The figure shows the relative SE of Top3 compared to the full RF, for different hyperparameter settings. Green bars: standard hyperparameters, as listed in Table 2. Red bars: hyperparameters that optimize the full RF. Blue bars: hyperparameters that optimize Top3. Hyperparameter tuning is based on SE.

hurts Top3, with a median relative SE of 1.32, compared to 1.22 in the standard case. By contrast, tuning on Top3 benefits Top3, with a median relative SE of 1.13. In a small minority of cases, tuning is not beneficial in terms of the relative loss.

Hence, despite its simplicity, Top3 can perform similar to the full RF given suitable hyperparameter choices. Overall, the relative performance of Top3 obtained under the standard setting is between the performance obtained under the two other hyperparameter settings. Furthermore, the full RF and Top3 require different sets of hyperparameters in order to perform well. Ideally, users of the Top k method should thus choose hyperparameters based on the performance of Top k itself, rather than the performance of the full RF.

4. Stylized Analytical Model

In this section, we construct a stylized analytical framework which helps explain our experimental findings presented in Section 3: For many datasets, simplified RFs perform similar to full RFs even for relatively small choices of k . This finding, and especially the possibility that simplified RFs may even outperform full RFs, deserves further investigation. Motivated by the structure of RF forecast distributions (see Section 2),

we consider a model in which the true forecast distribution is discrete with support points $\mathbf{u} = (u_i)_{i=1}^n$ and corresponding (true) probabilities $\omega^* = (\omega_i^*)_{i=1}^n$ that are positive and sum to one. Specifically, we let

$$\omega_i^* = \begin{cases} \theta^*/k & \text{if } i \in \mathcal{I} \\ (1 - \theta^*)/(n - k) & \text{if } i \notin \mathcal{I}, \end{cases} \quad (9)$$

where $\mathcal{I} \subseteq \{1, 2, \dots, n\}$ is a subset of “important” indexes with $|\mathcal{I}| = k$. The corresponding “important” probabilities $(\omega_i^*)_{i \in \mathcal{I}}$ sum to $\theta^* \in [0, 1]$, whereas the other, “unimportant”, probabilities sum to $1 - \theta^*$. To justify the notion of “important” probabilities, we will focus on choices of θ^* and k that satisfy $\theta^*/k > (1 - \theta^*)/(n - k)$.

In addition to the true forecast distribution just described, we consider an estimated forecast distribution that uses the same support points \mathbf{u} , together with possibly incorrectly estimated probabilities $\omega = (\omega_i)_{i=1}^n$. We assume that the estimated probabilities can be described by the following model:

$$\omega_i = \begin{cases} \theta Z_{1,i} & \text{if } i \in \mathcal{I} \\ (1 - \theta) Z_{2,i} & \text{if } i \notin \mathcal{I}, \end{cases} \quad (10)$$

where $\theta \in [0, 1]$, Z_1 is a draw from a Dirichlet distribution with k -dimensional parameter vector (d_1, \dots, d_1) , with $d_1 > 0$, such that each element of Z_1 has expected value $1/k$ and variance $(k - 1)/(k^2(kd_1 + 1))$. Similarly, Z_2 is a draw from another, independent Dirichlet distribution with $(n - k)$ -dimensional parameter vector (d_2, \dots, d_2) , where $d_2 > 0$. This means that the expected probabilities are given by

$$\mathbb{E}[\omega_i] = \begin{cases} \theta/k & \text{if } i \in \mathcal{I} \\ (1 - \theta)/(n - k) & \text{if } i \notin \mathcal{I}. \end{cases} \quad (11)$$

In the following, we assume that $2 \leq k \leq n - 2$, which ensures that there are at least two “important” and “unimportant” support points, respectively. This restriction ensures that weight estimation within both sets is a nontrivial problem.⁵

Thus, if $\theta \neq \theta^*$, the forecast model’s expected probabilities differ from the true ones in (9). The parameters d_1 and d_2 represent the precision of the forecast model’s probabilities around their expected values. Small values for d_1, d_2 indicate noisy probabilities, whereas large values for d_1, d_2 correspond to probabilities close to their expected values. This is a property of the variance of Dirichlet-distributed random variables as noted above for the case of Z_1 . Conceptually, the above model provides a stylized probabilistic description of the estimated probabilities ω produced by a forecasting method like RFs. Thereby, the model does not aim to specify the mechanism by which the forecasting method generates these probabilities.

Figure 5 illustrates this setup, with $n = 20$ and the important indexes given by $\mathcal{I} = \{1, 2, 3, 4, 5\}$. The left panel shows a situation in which the estimated probabilities are quite noisy ($d_1 = d_2 = 1$) but are correct in expectation ($\theta = \theta^*$). In the right panel, the estimated probabilities are less noisy ($d_1 = d_2 = 10$) but are false in expectation ($\theta = 0.4 \neq 0.8 = \theta^*$).

⁵If there was only one important support point, for example, the probability of this support point would necessarily be equal to θ , rendering weight estimation trivial.

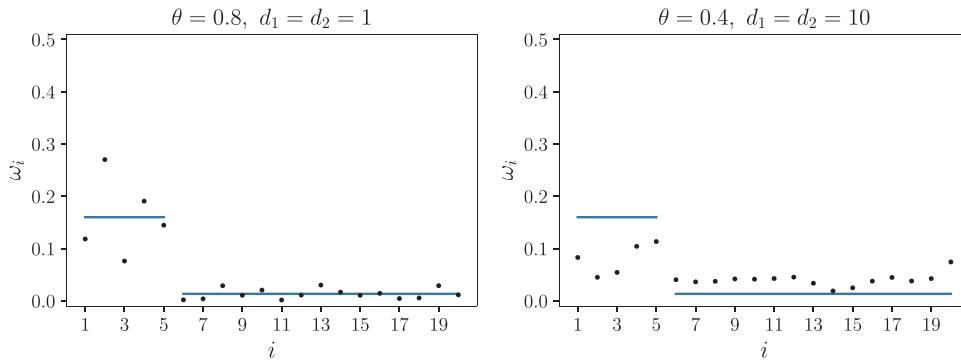


Figure 5. Simulated probability estimates. In both panels, we set $n = 20, k = 5$, and $\theta^* = 0.8$. The other parameters (θ, d_1 and d_2) are as indicated in the header. Horizontal segments represent true probabilities, dots represent estimated probabilities.

In the special case that $\theta = \theta^*, d_1 \rightarrow \infty, d_2 \rightarrow \infty$, the forecast model coincides with the true model. Furthermore, in the case $\theta = 1$, the forecast model is very similar to the “Top k ” strategy (retaining the k most important probabilities, rescaling them to sum to one, and setting all other probabilities to zero). The somewhat subtle difference between the analytical model and our practical implementation of Top k is that the indexes of the important weights are fixed in the analytical model (given by the set \mathcal{I}), whereas they are chosen as the k largest empirical weights in practice. Exact modeling of our practical procedure would seem to complicate the analysis substantially without necessarily yielding further insights. While stylized, the analytical model described above is flexible enough to cover various situations of applied interest. For example, the relation between “important” versus “unimportant” values of the true probabilities can be governed flexibly via the parameters n, k , and θ^* . While we assume that the set \mathcal{I} of important indexes is known to the forecast model, the possibility of a poor forecasting model can be represented by a value θ that differs substantially from θ^* , and/or small values of d_1, d_2 that correspond to noisy estimates. Thus, the forecast model could even yield estimates of the “unimportant” probabilities that greatly exceed those of the “important” ones.⁶ For given support points \mathbf{u} and estimated probabilities ω , the expected squared error and expected CRPS implied by the analytical framework are given by

$$\mathbb{E}[\text{SE}(\omega, \mathbf{u})] = \sum_{i=1}^n \omega_i^* (u_i - \sum_{j=1}^n \omega_j u_j)^2, \quad (12)$$

$$\mathbb{E}[\text{CRPS}(\omega, \mathbf{u})] = \sum_{i=1}^n \sum_{j=1}^n \omega_i (\omega_j^* - \omega_j/2) |u_i - u_j|; \quad (13)$$

the expression for the CRPS follows from adapting the representation in eq. 2 of Jordan, Krüger, and Lerch (2019). In both cases, the expected value is computed with respect to the discrete distribution with support points \mathbf{u} and associated true probabilities ω^* . As noted in (10), we cast the predicted probabilities ω as scaled draws from two Dirichlet distributions. We further assume that the support points \mathbf{u} are n draws from a standard normal distribution; these are mutually independent and independent of ω . Using these assumptions, we obtain the following

⁶In practice, where the set of important indexes \mathcal{I} is not known, this situation corresponds to one in which the largest empirical weights are not helpful for predicting new test sample cases.

expressions for the (unconditionally) expected squared error and CRPS:

$$\begin{aligned} \mathbb{E}[\text{SE}] &= \int \int \mathbb{E}[\text{SE}(\omega, \mathbf{u})] dF_\omega(\omega) dF_{\mathbf{u}}(\mathbf{u}) \\ &= 1 - 2 \left\{ \frac{\theta^* \theta}{k} + \frac{(1 - \theta^*)(1 - \theta)}{n - k} \right\} \\ &\quad + \frac{\theta^2}{k} + \frac{(1 - \theta)^2}{n - k} + \frac{\theta^2(k - 1)}{k(d_1 k + 1)} \\ &\quad + \frac{(1 - \theta)^2(n - k - 1)}{(n - k)(d_2(n - k) + 1)}, \end{aligned} \quad (14)$$

$$\begin{aligned} \mathbb{E}[\text{CRPS}] &= \int \int \mathbb{E}[\text{CRPS}(\omega, \mathbf{u})] dF_\omega(\omega) dF_{\mathbf{u}}(\mathbf{u}) \\ &= \frac{1}{\sqrt{\pi}} \mathbb{E}[\text{SE}], \end{aligned} \quad (15)$$

where F_ω is the distribution of the estimated probabilities that is implied by our model setup, and $F_{\mathbf{u}}$ is the joint distribution of n independent standard normal variables. The proof can be found in Section C of the online supplement. The result that the expressions for $\mathbb{E}[\text{SE}]$ and $\mathbb{E}[\text{CRPS}]$ are identical up to a factor of $\sqrt{\pi}$ is a somewhat idiosyncratic implication of our model setup.

In order to interpret the implications of these formulas, we compare a forecasting method with $\theta < 1$ (representing standard RFs) to a method with $\theta = 1$ (representing Top k) in the following.

For given values of n, k , and θ^* , both $\mathbb{E}[\text{SE}]$ and $\mathbb{E}[\text{CRPS}]$ attain their theoretical minimum at $\theta = \theta^*, d_1 \rightarrow \infty$ and $d_2 \rightarrow \infty$.⁷ This result is unsurprising: Under the stated conditions, the forecast model coincides with the true model, that is, $\omega = \omega^*$ with probability one. Since the squared error is strictly consistent for the mean (and, similarly, the CRPS is a strictly proper scoring rule), the true model must yield the smallest possible expected score.⁸ As both expected score functions are continuous in θ, d_1 and d_2 , this implies that if θ is sufficiently close to θ^* , and

⁷Proof: $\frac{\partial \mathbb{E}[\text{SE}]}{\partial d_i} < 0$ for $i = 1, 2$; this holds for all values of $\theta, \theta^*, n, k, d_1$, and d_2 . It is hence optimal to let d_1, d_2 go to infinity. Next consider the limiting expression of $\mathbb{E}[\text{SE}]$ as $d_1, d_2 \rightarrow \infty$. Minimizing this expression with respect to θ yields the solution $\theta = \theta^*$.

⁸While the possibility of exactly matching the true model is unrealistic in practice, the requirement that the true model perform best is conceptually plausible, and is the main idea behind forecast evaluation via proper scoring rules and related tools.

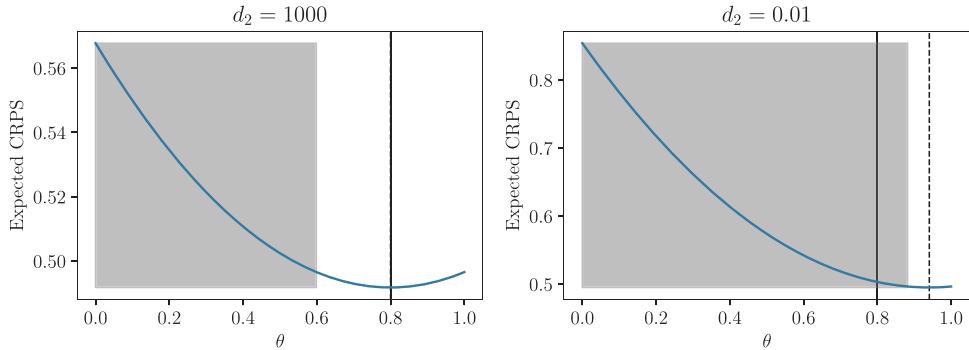


Figure 6. Expected CRPS as a function of θ . The left panel refers to $d_2 = 1000$ (i.e., precise estimates of “unimportant” probabilities), whereas the right panel assumes $d_2 = 0.01$ (i.e., noisy estimates). The other parameters are set as follows: $n = 100, k = 5, \theta^* = 0.8, d_1 = 1000$. Solid vertical line marks θ^* , dashed vertical line marks best value for θ . Shaded area marks range of values for θ that perform worse than $\theta = 1$ (corresponding to the Topk method).

d_1, d_2 are sufficiently large, then the standard approach will outperform the Topk method.

Conversely, the following conditions favor the Topk method over the standard approach:

- $|\theta - \theta^*| > |1 - \theta^*|$, that is the standard method’s implicit assumption that $\theta^* = \theta$ is worse than Topk’s implicit assumption that $\theta^* = 1$
- d_1 is large, that is important probabilities are estimated precisely
- d_2 is small, that is estimates of unimportant probabilities are noisy

If these conditions, or an appropriate combination thereof, hold, then the Topk approach can be expected to perform well.

Figure 6 illustrates the above discussion. In the left panel (with $d_2 = 1000$), the “unimportant” probabilities are estimated very precisely. Here the Topk method is superior only to values $\theta \leq 0.6$ that are clearly smaller than the true parameter $\theta^* = 0.8$. In the right panel ($d_2 = 0.01$), the estimates of the unimportant probabilities are very noisy. Hence it is beneficial to focus on the important probabilities which are estimated precisely (since $d_1 = 1000$). Accordingly, the Topk method—which focuses on the important probabilities exclusively—is superior to a wide range of values for θ . Interestingly, this range includes the true parameter $\theta = \theta^*$, that is, the Topk method can be beneficial even if the probability estimates are correct in expectation.

5. Discussion

This article has considered simplified RF forecast distributions that consist of a small number k of support points, in contrast to thousands of support points (possibly equal to n , the size of the training set) of the original forecast distribution. The Topk forecast distribution can be viewed as a collection of k scenarios with attached probabilities. It hence simplifies communication and improves interpretability of the probabilistic forecast. Our empirical results in Tables 2 and 4 imply that simplified distributions using 5 or 10 support points often attain similar performance as the original forecast distribution, while larger choices of k , for example, 20 or 50, even increase performance slightly in many cases. Our analytical framework in Section 4 offers a theoretical rationale for these results. In particular, if

small weights are estimated imprecisely by the full RF, using only the k largest weights can improve performance. Our empirical analysis further shows that when tuning hyperparameters to the target value for k , even $k = 3$ can yield very good results.

A limitation of our proposed method is that focusing on a small number of weights is not promising if the original RF weights are close to equal. This occurs, in particular, if the regressors are entirely uninformative so that the predictive distribution for a new test point $X = \mathbf{x}_0$ is given by the equally weighted unconditional distribution of the training responses $(y_i)_{i=1}^n$, independently of \mathbf{x}_0 . The latter distribution is consistent for the true forecast distribution as $n \rightarrow \infty$, while focusing on a finite number of k weights prevents consistency. In terms of predictability, two of the 18 datasets we study are reminiscent of this situation (see Section 3.1).

While we have focused on the tradeoff between simplicity (as measured by k) and statistical forecasting performance, the optimal choice of k depends on the preferences of the forecaster and the forecast users. These preferences are necessarily subjective and likely application specific. For choosing k in practice, we recommend to first assess the statistical performance of various choices of k for the application or dataset of interest, similar to our empirical analysis in Section 3. If ease of communication is not a concern, forecasters can simply pick the value of k that optimizes (estimated) performance. Conversely, if ease of communication matters, smaller values for k may be preferable even at the price of reduced statistical performance. To navigate this tradeoff in practice, forecasters might interview potential users about their perceived cognitive costs of various choices of k . For example, Altig et al. (2022) argue that $k = 5$ resonates well with participants of an online survey on firm performance.

Finally, the idea of simplifying predictions by sparsifying observation weights can possibly be extended to methods other than RFs. For example, the predictions of linear regression, ridge regression, and kernel representations of certain neural network models can be represented as weighted sums of training sample outcomes (see e.g., Jacot, Gabriel, and Hongler 2018; Goulet Coulombe, Göbel, and Klieber 2024).

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

We thank conference participants at the International Symposium on Forecasting (Dijon 2024) as well as Andreas Eberl and two anonymous reviewers for helpful comments.

References

Abbas, A. E., and Howard, R. A. (2015), *Foundations of Decision Analysis*, Boston: Pearson. [1]

Altig, D., Barrero, J. M., Bloom, N., Davis, S. J., Meyer, B., and Parker, N. (2022), “Surveying Business Uncertainty,” *Journal of Econometrics*, 231, 282–303. [1,10]

Beck, E., Kozbur, D., and Wolf, M. (2023), “Hedging Forecast Combinations with an Application to the Random Forest,” preprint, arXiv:2308.15384. [3]

Biau, G., and Scornet, E. (2016), “A Random Forest Guided Tour,” *TEST*, 25, 197–227. [3]

Breiman, L. (2001), “Random Forests,” *Machine Learning*, 45, 5–32. [1]

Cevid, D., Michel, L., Naf, J., Bühlmann, P., and Meinshausen, N. (2022), “Distributional Random Forests: Heterogeneity Adjustment and Multivariate Distributional Regression,” *Journal of Machine Learning Research*, 23, 1–79. [3,4]

Cramer, E. Y., Ray, E. L., Lopez, V. K., Bracher, J., et al. (2022), “Evaluation of Individual and Ensemble Probabilistic Forecasts of COVID-19 Mortality in the United States,” *Proceedings of the National Academy of Sciences*, 119, e2113561119. [1]

DIW. (2022), “SOEP-Übungsdatensatz, Daten der Jahre 2015–2019,” Data set. DOI:10.5684/soep.practice.v36 [2]

Friedman, J. H., and Popescu, B. E. (2008), “Predictive Learning via Rule Ensembles,” *Annals of Applied Statistics*, 2, 916–954. [3]

Gneiting, T., and Katzfuss, M. (2014), “Probabilistic Forecasting,” *Annual Review of Statistics and Its Application*, 1, 125–151. [4]

Gneiting, T., and Raftery, A. E. (2007), “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association*, 102, 359–378. [4]

Gneiting, T., and Resin, J. (2023), “Regression Diagnostics Meets Forecast Evaluation: Conditional Calibration, Reliability Diagrams, and Coefficient of Determination,” *Electronic Journal of Statistics*, 17, 3226–3286. [6]

Goebel, J., Grabka, M. M., Liebig, S., Kroh, M., Richter, D., Schröder, C., and Schupp, J. (2019), “The German Socio-Economic Panel (SOEP),” *Jahrbücher für Nationalökonomie und Statistik*, 239, 345–360. [2]

Goulet Coulombe, P., Göbel, M., and Klieber, K. (2024), “Dual Interpretation of Machine Learning Forecasts,” preprint, arXiv:2412.13076. [10]

Grinsztajn, L., Oyallon, E., and Varoquaux, G. (2022), “Why Do Tree-based Models Still Outperform Deep Learning on Typical Tabular Data?” in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, pp. 507–520. [3,5,6]

Haddouchi, M., and Berrado, A. (2019), “A Survey of Methods and Tools Used for Interpreting Random Forest,” in *2019 1st International Conference on Smart Systems and Data Science (ICSSD)*, pp. 1–6. [3]

Hastie, T., Tibshirani, R., and Friedman, J. (2009), *The Elements of Statistical Learning Data Mining, Inference, and Prediction* (2nd ed.), New York: Springer. [3]

Jacot, A., Gabriel, F., and Hongler, C. (2018), “Neural Tangent Kernel: Convergence and Generalization in Neural Networks,” in *Advances in Neural Information Processing Systems* (Vol. 31), eds. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Curran Associates, Inc. [10]

Jordan, A. I. (2016), “Facets of Forecast Evaluation,” PhD thesis, Karlsruher Institut für Technologie. [5]

Jordan, A. I., Krüger, F., and Lerch, S. (2019), “Evaluating Probabilistic Forecasts with ScoringRules,” *Journal of Statistical Software*, 90, 1–37. [5,9]

Krüger, F., Clark, T. E., and Ravazzolo, F. (2017), “Using Entropic Tilting to Combine BVAR Forecasts with External Nowcasts,” *Journal of Business & Economic Statistics*, 35, 470–485. [1]

Laio, F., and Tamea, S. (2007), “Verification Tools for Probabilistic Forecasts of Continuous Hydrological Variables,” *Hydrology and Earth System Sciences*, 11, 1267–1277. [5]

Liaw, A., and Wiener, M. (2002), “Classification and Regression by Random Forest,” *R News*, 2, 18–22. [7]

Lin, Y., and Jeon, Y. (2006), “Random Forests and Adaptive Nearest Neighbors,” *Journal of the American Statistical Association*, 101, 578–590. [1,3,4,7]

Lundberg, S. M., and Lee, S.-I. (2017), “A Unified Approach to Interpreting Model Predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, pp. 4768–4777. [3]

Matheson, J. E., and Winkler, R. L. (1976), “Scoring Rules for Continuous Probability Distributions,” *Management Science*, 22, 1087–1096. [4]

Meinshausen, N. (2006), “Quantile Regression Forests,” *Journal of Machine Learning Research*, 7, 983–999. [1,3,4]

——— (2010), “Node Harvest,” *Annals of Applied Statistics*, 4, 2049–2072. [3]

Nan, F., Wang, J., and Saligrama, V. (2016), “Pruning Random Forests for Prediction on a Budget,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS’16, pp. 2342–2350. [3]

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011), “Scikit-Learn: Machine Learning in Python,” *Journal of Machine Learning Research*, 12, 2825–2830. [3,5]

Probst, P., Wright, M. N., and Boulesteix, A.-L. (2019), “Hyperparameters and Tuning Strategies for Random Forest,” *WIREs Data Mining and Knowledge Discovery*, 9, e1301. [4,7]

R Core Team. (2022), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing. Available at <https://www.R-project.org/> (last accessed: February 4, 2024). [5]

Raftery, A. E. (2016), “Use and Communication of Probabilistic Forecasts,” *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 9, 397–410. [1]

Rasp, S., and Lerch, S. (2018), “Neural Networks for Postprocessing Ensemble Weather Forecasts,” *Monthly Weather Review*, 146, 3885–3900. [1]

Székely, G. J., and Rizzo, M. L. (2017), “The Energy of Data,” *Annual Review of Statistics and Its Application*, 4, 447–479. [4]

Taieb, S. B., Taylor, J. W., and Hyndman, R. J. (2021), “Hierarchical Probabilistic Forecasting of Electricity Demand with Smart Meter Data,” *Journal of the American Statistical Association*, 116, 27–43. [1]

van Rossum, G., et al. (2011), “Python Programming Language,” available at <https://www.python.org> (last accessed: February 4, 2024). [5]

Wager, S., and Athey, S. (2018), “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests,” *Journal of the American Statistical Association*, 113, 1228–1242. [3]

Winkler, R. L. (1996), “Scoring Rules and the Evaluation of Probabilities,” *TEST*, 5, 1–26. [4]

Wright, M. N., and Ziegler, A. (2017), “ranger: A Fast Implementation of Random Forests for High Dimensional Data in C++ and R,” *Journal of Statistical Software*, 77, 1–17. [5,7]

Zhao, X., Wu, Y., Lee, D. L., and Cui, W. (2019), “IForest: Interpreting Random Forests via Visual Analytics,” *IEEE Transactions on Visualization and Computer Graphics*, 25, 407–416. [3]