



HAL
open science

Quantile Super Learning for independent and online settings with application to solar power forecasting

Herbert Susmann, Antoine Chambaz

► **To cite this version:**

Herbert Susmann, Antoine Chambaz. Quantile Super Learning for independent and online settings with application to solar power forecasting. 2023. hal-04262836

HAL Id: hal-04262836

<https://hal.science/hal-04262836>

Preprint submitted on 27 Oct 2023

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

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

Copyright

Quantile Super Learning for independent and online settings with application to solar power forecasting

Herbert Susmann^{a,1,*}, Antoine Chambaz^b

^a*CEREMADE (UMR 7534), Université Paris-Dauphine PSL, Place du Maréchal de Lattre de Tassigny, Paris, 75016, France*

^b*MAP5 (UMR 8145), Université Paris Cité, 75006, Paris, France*

Abstract

Estimating quantiles of an outcome conditional on covariates is of fundamental interest in statistics with broad application in probabilistic prediction and forecasting. We propose an ensemble method for conditional quantile estimation, Quantile Super Learning, that combines predictions from multiple candidate algorithms based on their empirical performance measured with respect to a cross-validated empirical risk of the quantile loss function. We present theoretical guarantees for both i.i.d. and online data scenarios. The performance of our approach for quantile estimation and in forming prediction intervals is tested in simulation studies. Two case studies related to solar energy are used to illustrate Quantile Super Learning: in an i.i.d. setting, we predict the physical properties of perovskite materials for photovoltaic cells, and in an online setting we forecast ground solar irradiance based on output from dynamic weather ensemble models.

Keywords: cross validation, online learning, quantile regression

1. Introduction

Estimating the quantiles of an outcome conditional on covariates is a foundational task in statistics. Many algorithms have been developed for that purpose, including versions of linear regression, generalized additive models, random forests, and neural networks, to name only a few (Koenker, 2005; Athey et al., 2019; Cannon, 2011; Fasiolo et al., 2020). On any particular dataset, however, it is almost never known a-priori which method will perform best. Ensemble algorithms combine the predictions of multiple candidate algorithms according to their empirical performance, obviating the need to choose between the available algorithms in advance. In particular, Super Learning combines the predictions of candidate algorithms according to their cross-validated risk with respect to a loss function chosen on a case by case basis depending on the task at hand (van der Laan et al., 2007).

In this work we present the Quantile Super Learner (QSL), a novel ensemble learning algorithm tailored to estimating conditional quantiles. To illustrate the main ideas, suppose we have n independent and identically distributed (i.i.d.) observations (O_1, \dots, O_n) of a

*Corresponding author.

Email address: `herbert.susmann@dauphine.psl.eu` (Herbert Susmann)

generic variable $O = (X, Y)$ drawn from a law P_0 , with $X \in \mathcal{X}$ a set of covariates and $Y \in \mathbb{R}$ a univariate outcome. Our goal is to estimate the α -conditional quantile of Y given X , denoted $\psi_{P_0}^\alpha(X)$ where $\psi_{P_0}^\alpha \in \Psi$ with Ψ the set of all functionals mapping $\mathcal{X} \rightarrow \mathbb{R}$. At our disposal are a set of algorithms $\hat{\psi}_k^\alpha$, $k = 1, \dots, K$ that map a dataset to an estimator of the α -conditional quantile. No assumptions are necessary about how these algorithms work, and in our simulations and case studies we use a variety methods including quantile regression, gradient boosting machines, and neural networks, among others.

As we are not likely to know a-priori which algorithm will perform best for a particular dataset, we choose between them by evaluating their performance with respect to the quantile loss function L^α , given for any $\psi \in \Psi$ by

$$(x, y) \mapsto L^\alpha(\psi)(x, y) = \begin{cases} \alpha|y - \psi(x)|, & \text{if } y > \psi(x) \\ (1 - \alpha)|y - \psi(x)|, & \text{if } y \leq \psi(x). \end{cases}$$

We choose this loss function because its expected value under P_0 is minimized by the true conditional quantile functional:

$$\psi_{P_0}^\alpha \in \arg \min_{\psi \in \Psi} \mathbb{E}_{P_0}[L^\alpha(\psi)(X, Y)] = \arg \min_{\psi \in \Psi} R_{P_0}^\alpha(\psi),$$

where we have defined $R_{P_0}^\alpha(\psi)$ to be the *risk* of the algorithm ψ with respect to P_0 . If we had oracular knowledge of the true data generating distribution P_0 we could evaluate the true risk $R_{P_0}^\alpha$ directly for the output of each algorithm and choose the one with the lowest risk. In practice we do not have access to P_0 , so we must approximate the true risk. The strategy used in Super Learning is to use a cross-validated risk as an estimator of the true risk, and then select the algorithm that minimizes this cross-validated risk. Our main theoretical results are in the form of a bound on the difference between the risk of the algorithm selected using the cross-validated risk and the risk of the algorithm that minimizes the true risk, following the strategy of (Wu and Benkeser, 2022). Notably, these results are established without any regularity assumptions on the data generating distribution.

So far, we have discussed the case where the data are i.i.d. draws from a probability law. However, in many scenarios i.i.d. assumptions are not justified, such as for time-series data. The key difference between Super Learning in the i.i.d. and sequential cases is in the cross-validation scheme used to calculate the empirical risk of the candidate algorithms. In the i.i.d. case, V -fold cross-validation is typically used, in which the dataset is split uniformly at random into V couples of training and test sets. For the online setting, the collection of all previously observed data is used as a training set and the next observation (or set of observations) is used as a test set. Our main theoretical result for the online setting is similar to that of the i.i.d. result, providing bounds on the excess risk of the algorithm chosen based on minimizing the online cross-validated risk. However, for the online setting it is necessary to introduce mild regularity assumptions on the data-generating process. Using ideas from Steinwart and Christmann (2011), we make a margin assumption concerning the behavior of the conditional law around the quantile being estimated, which allows us to establish excess risk bounds.

One of the potential use cases of the quantile Super Learner in both i.i.d. and online settings is to form prediction intervals by separately estimating a lower and upper quantile.

Intuitively, we would expect that if we can do a good job estimating each of these quantiles, then the prediction intervals will also perform well. However, there is no theoretical guarantee that such quantile estimates will yield prediction intervals with desired frequency characteristics in finite samples. We investigate this use case empirically in simulations and case studies.

Prior Work. Early proposals of combining predictions from multiple base learners include stacked generalization, described by Wolpert (1992) and Breiman (1996). The theoretical foundations of the approach were formalized by van der Laan and Dudoit (2003), van der Laan et al. (2006), and van der Vaart et al. (2006); the name “Super Learner” was subsequently coined by van der Laan et al. (2007). Early development of Super Learning primarily focused on estimating conditional means using a squared-error loss function in the context of i.i.d. data. Since then, extensions to other loss functions have included the Area Under the Curve (AUC) loss function (LeDell et al., 2016) and the Huber loss function (Wu and Benkeser, 2022), among others. Theory for online (sequential) ensemble learning within the Super Learning framework that incorporates the statistical dependence of the data was developed in Benkeser et al. (2018) and Ecoto et al. (2021). Fakoor et al. (2023) provide a comprehensive account of ensemble methods for quantile estimation, including detailed empirical comparisons. Our work is in a similar vein, but provides formal theoretical guarantees for the proposed Super Learning based approach. Sun et al. (2023) propose an online estimator for quantiles within a parametric model and using a smooth approximation of the quantile loss function. In this framework, they show that their “renewable estimator” is consistent, asymptotically normal, and that it enjoys an oracle property. In contrast, our approach is fully non-parametric.

Outline. The rest of the paper unfolds as follows. In Section 2 we develop the QSL for i.i.d. data and establish an oracle inequality showing that the estimator is asymptotically equivalent to the best performing candidate algorithm. In Section 3 we extend the QSL to the setting of sequential data and establish similar oracle inequalities as in the i.i.d. case. In Section 4 we investigate the finite-sample performance of the QSL in simulations. In Section 5 we present two case studies related to solar energy: predicting the physical properties of perovskite materials for photovoltaic applications and forecasting solar irradiance based on the output of deterministic numerical weather prediction models.

2. Independent setting

Our initial results develop a QSL in a setting where the observed data represent i.i.d. draws from an underlying distribution. Let (O_1, \dots, O_n) be n i.i.d. observations of a generic variable $O = (X, Y)$ from the law P_0 on $\mathcal{O} = \mathcal{X} \times \mathbb{R}$, where $X \in \mathcal{X}$ is a set of covariates and $Y \in \mathbb{R}$ is a univariate outcome. We assume that P_0 falls in statistical model \mathcal{M} . For all $P \in \mathcal{M}$ and $\alpha \in (0, 1)$, define the α -quantile of Y conditional on X as the possibly set-valued functional

$$\psi_P^\alpha(X) := \{y \in \mathbb{R} : P(Y < y | X) \geq \alpha, P(Y \geq y | X) \geq 1 - \alpha\}.$$

Our goal is to estimate $x \mapsto \psi_{P_0}^\alpha(x)$ for a given α , the conditional α -quantile under the true data generating distribution P_0 . For the i.i.d. setting, we make two key simplifying assumptions:

Assumption A1 (Unique Quantiles). *The quantile function $\psi_{P_0}^\alpha(x)$ is a singleton for P_X -almost all $x \in X$.*

Assumption A2 (Outcome Boundedness). *It holds P_0 almost surely that $|Y| \leq C_0 < \infty$. Note that in practice, the value of C_0 does not need to be known.*

Next, we establish that ψ_P^α is a minimizer of a particular loss function. Let Ψ be the set of all (measurable) functions mapping \mathcal{X} to \mathbb{R} . The quantile loss function L^α for the α -quantile is given by

$$(x, y) \mapsto L^\alpha(\psi)(x, y) := \begin{cases} \alpha|y - \psi(x)|, & \text{if } y > \psi(x) \\ (1 - \alpha)|y - \psi(x)|, & \text{if } y \leq \psi(x) \end{cases}$$

for any $(x, y) \in \mathcal{O}$ and $\psi \in \Psi$. The *risk* of ψ under the loss function L^α relative to a distribution $P \in \mathcal{M}$ is defined as

$$R_P^\alpha(\psi) := \mathbb{E}_P[L^\alpha(\psi)(O)] = PL^\alpha(\psi), \quad (1)$$

where we use the notation $Pf = \mathbb{E}_P[f(O)] = \int fdP$. Note that the true conditional quantile function minimizes the quantile loss:

$$\psi_{P_0}^\alpha = \arg \min_{\psi \in \Psi} R_{P_0}^\alpha(\psi).$$

This well-known fact serves as the basis of quantile regression (Gneiting, 2011; Koenker, 2005).

2.1. Super Learning

An *algorithm* to learn $\psi_{P_0}^\alpha$ is a function mapping any finite set $\{o_1, \dots, o_M\}$ of M elements of \mathcal{O} , viewed as the measure $M^{-1} \sum_{m=1}^M \text{Dirac}(o_m)$, to an element of Ψ . Suppose we have K such algorithms $\hat{\psi}_1^\alpha, \dots, \hat{\psi}_K^\alpha$ which seek to learn $\psi_{P_0}^\alpha$. Super Learning amounts to identifying which of the algorithms performs best.

Discrete Super Learner.. The discrete Super Learner identifies the best performing algorithm among the candidate algorithms as measured by their cross-validated risks. To formalize the cross-validation scheme, we introduce $B_n \in \{0, 1\}^n$, a random vector drawn independently of O_1, \dots, O_n such that $\sum_{i=1}^n B_n(i) \approx np$ for some user-supplied proportion p . The observation O_i falls in the training set if $B_n(i) = 0$, and in the testing set if $B_n(i) = 1$. The B_n -specific training and testing datasets are represented by the empirical distributions P_{n, B_n}^0 and P_{n, B_n}^1 . For instance, to implement V -fold cross-validation, we draw B_n from the uniform distribution on $\{b_1, \dots, b_V\} \subset \{0, 1\}^n$ where each b_v satisfies $\sum_{i=1}^n b_v(i) \approx n/V$ (a proportion $(V-1)/V$ of data are used for training, and the rest for testing) and, for every $1 \leq i \leq n$, $\sum_{v=1}^V b_v(i) = 1$ (each observation is used once for testing).

The oracle cross-validated risk of algorithm $\widehat{\psi}_k^\alpha$ with respect to a distribution $P \in \mathcal{M}$ is then defined by

$$\widetilde{R}_{n,P}^\alpha(\widehat{\psi}_k^\alpha) := \mathbb{E}_{B_n} \left[PL^\alpha(\widehat{\psi}_k^\alpha(P_{n,B_n}^0)) \right]. \quad (2)$$

The empirical version of the oracle cross-validated risk is simply obtained by substituting P_{n,B_n}^1 for P in (2):

$$\widehat{R}_n^\alpha(\widehat{\psi}_k^\alpha) := \mathbb{E}_{B_n} \left[P_{n,B_n}^1 L^\alpha(\widehat{\psi}_k^\alpha(P_{n,B_n}^0)) \right].$$

The discrete Super Learner selector minimizes the empirical cross-validated risk:

$$\widehat{\kappa}_n := \arg \min_{k \in \llbracket K \rrbracket} \widehat{R}_n^\alpha(\widehat{\psi}_k^\alpha),$$

using the notation $\llbracket K \rrbracket := \{1, 2, \dots, K\}$. The corresponding algorithm $\widehat{\psi}_{\widehat{\kappa}_n}^\alpha$ is referred to as the discrete Super Learner.

Continuous Super Learner. The continuous Super Learner considers a richer class of algorithms taking the form of convex combinations of the original candidate algorithms. Let Π be the K -simplex: that is, the set of $\pi \in (\mathbb{R}_+)^K$ such that $\sum_{k=1}^K \pi_k = 1$. The new generic candidate algorithms take the form

$$\widehat{\psi}_\pi^\alpha := \sum_{k=1}^K \pi_k \widehat{\psi}_k^\alpha \quad (3)$$

for any $\pi \in \Pi$. Let Π_n be a finite subset of Π such that the cardinality of Π_n grows at most polynomially with n . The continuous Super Learner is found by finding the weights $\pi \in \Pi_n$ that minimize the empirical cross-validated risk:

$$\widehat{\pi}_n \in \arg \min_{\pi \in \Pi_n} \widehat{R}_n^\alpha(\widehat{\psi}_\pi^\alpha).$$

The algorithm $\widehat{\psi}_{\widehat{\pi}_n}^\alpha$ is referred to as the continuous Super Learner. Note that the continuous Super Learner is simply the discrete Super Learner when the collection of candidate algorithms is $\{\widehat{\psi}_\pi^\alpha : \pi \in \Pi_n\}$. As such, we focus on analyzing the properties of the discrete Super Learner, as the results carry over to the continuous Super Learner.

2.2. Oracle Inequalities

We compare the discrete Super Learner against an oracle selector which identifies the candidate algorithm that has the best oracle cross-validated risk with respect to the law P_0 :

$$\widetilde{\kappa}_n := \arg \min_{k \in \llbracket K \rrbracket} \widetilde{R}_{n,P_0}^\alpha(\widehat{\psi}_k^\alpha).$$

The corresponding algorithm $\widehat{\psi}_{\widetilde{\kappa}_n}^\alpha$ is referred to as the oracle Super Learner.

Let $\tilde{\psi}_{P_0}^\alpha$ denote the oracle algorithm that constantly outputs the true conditional α -quantile $\psi_{P_0}^\alpha$. Our theoretical results compare the excess risk of the discrete Super Learner to the excess risk of the oracle Super Learner, that is

$$\tilde{R}_{n,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_n}^\alpha) - \tilde{R}_{n,P_0}^\alpha(\tilde{\psi}_{P_0}^\alpha) \text{ vs. } \tilde{R}_{n,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_n}^\alpha) - \tilde{R}_{n,P_0}^\alpha(\tilde{\psi}_{P_0}^\alpha),$$

where

$$\hat{R}_{n,P_0}^\alpha(\tilde{\psi}_{P_0}^\alpha) := \mathbb{E}_{B_n} [P_{n,B_n}^1 L^\alpha(\psi_{P_0}^\alpha)].$$

We are now ready to state the main result of this section, an excess risk bound for the discrete Super Learner.

Theorem 1 (Excess risk bounds for discrete Super Learner). *Assume that the number of candidate algorithms grows at most polynomially in n : that is, $K = O(n^a)$ for some $a > 0$. Also assume that, for all $k \in \llbracket K \rrbracket$, $\hat{\psi}_k^\alpha$ only outputs functions which are uniformly bounded by C_0 . Then*

$$\mathbb{E}_{P_0} \left[\tilde{R}_{n,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_n}^\alpha) - \tilde{R}_{n,P_0}^\alpha(\tilde{\psi}_{P_0}^\alpha) \right] \leq \mathbb{E}_{P_0} \left[\hat{R}_{n,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_n}^\alpha) - \hat{R}_{n,P_0}^\alpha(\tilde{\psi}_{P_0}^\alpha) \right] + O\left(\frac{\log(n)}{n^{1/2}}\right).$$

3. Sequential setting

In the sequential setting we gain access to the observations in batches. In particular, we are interested in the setting where, at each time point, we gain access to a new batch of observations, one for each element in an index set \mathcal{J} . For example, we may have multiple time series corresponding to several locations, where each location generates a new observation at each timepoint. Note that if $|\mathcal{J}| = 1$, then the problem reduces to the case where we observe a single time series.

Let $(\bar{O}_t)_{t \geq 1}$ be a time-ordered sequence where t indexes time. Each \bar{O}_t is the set of observations $\bar{O}_t = (O_{j,t} : j \in \mathcal{J})$. The observations then decompose as $O_{j,t} = (X_{j,t}, Y_{j,t})$ for $j \in \mathcal{J}, t \geq 1$. Let P_0 be the joint law of the observed data, which we assume falls in a statistical model \mathcal{M} . For all $t \geq 2$, introduce the σ -field $F_{t-1} := \sigma(O_{j,\tau} : j \in \mathcal{J}, 1 \leq \tau < t)$ generated by past observations (with $F_0 := \emptyset$ by convention). For all $P \in \mathcal{M}$ and $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$ define the α -quantile of Y conditional on X at location j and time τ as the possibly set-valued

$$\psi_{P,j,t}^\alpha(X_{j,t}) := \{y \in \mathbb{R} : P(Y_{j,t} < y \mid X_{j,t}) \geq \alpha, P(Y_{j,t} \geq y \mid X_{j,t}) \geq 1 - \alpha\},$$

where $\alpha \in (0, 1)$ is taken as fixed. The sequential setting requires several more assumptions in addition to those adopted in the i.i.d. case. First, we make the simplifying assumption of common support of the covariates across all locations and time points.

Assumption B1 (Common Support). *For all $P \in \mathcal{M}$ there exists \mathcal{S} such that, for all $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$,*

$$\text{Supp}(P_{X_{j,\tau}}) = \mathcal{S},$$

where $P_{X_{j,\tau}}$ is the marginal law of X under P at location j and time τ .

As in the i.i.d. setting, we will also assume the existence of unique quantiles and that the outcomes are uniformly bounded.

Assumption B2 (Unique Quantiles). *For all $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$, the α -quantile $\psi_{P_0}^\alpha(x)$ is a singleton for all $x \in \mathcal{S}$.*

Assumption B3 (Outcome Boundedness). *It holds P_0 -almost surely that, for all $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$, $|Y_{j,t}| \leq C_0 < \infty$.*

We also make a Markovian assumption that all information about the outcome at a particular location and time point is encoded in prior observations and the covariates for that location and time point.

Assumption B4 (Markov). *For every $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$, it holds P_0 -almost surely that*

$$P_0(Y_{j,\tau} | (X_{j,\tau} : j \in \mathcal{J})) = P_0(Y_{j,\tau} | X_{j,\tau}).$$

Finally, the next assumption guarantees that we can learn the quantile function based on the observed time series $\bar{O}_1, \dots, \bar{O}_t$.

Assumption B5 (Stationarity). *There exists $\Psi_{P_0}^\alpha \in \Psi$ such that, for every $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$,*

$$\arg \min_{\psi \in \Psi} \mathbb{E}_{P_0}[L^\alpha(\psi)(O_{j,\tau})] = \psi_{P_0}^\alpha.$$

(Informally, the assumption states that $\Psi_{P_0, j, \tau}^\alpha = \Psi_{P_0}^\alpha$ for all $(j, \tau) \in \mathcal{J} \times \mathbb{N}^$.)*

We prove the oracle inequalities for Super Learning in the sequential setting by making a regularity assumption on the P_0 -conditional laws given F_t . In order to state the assumption, we need the following two definitions drawn from (Steinwart and Christmann, 2011).

Definition 1 (Quantiles of type q (Steinwart and Christmann, 2011)). *Let Q be a distribution with $\text{Supp}(Q) \subset [-1, 1]$. Set arbitrarily $\alpha \in (0, 1)$, let $\psi_Q^\alpha := \{t \in \mathbb{R} : Q((-\infty, t]) \geq \alpha, Q([t, \infty)) \geq 1 - \alpha\}$ be the α -quantile of Q , and assume that ψ_Q^α is a singleton. The distribution Q is said to have an α -quantile of type $q \in (1, \infty)$ if there exist constants $\alpha_Q \in (0, 2]$ and $b_Q > 0$ such that*

$$\begin{aligned} Q((\psi_Q^\alpha - s, \psi_Q^\alpha)) &\geq b_Q s^{q-1}, \\ Q((\psi_Q^\alpha, \psi_Q^\alpha + s)) &\geq b_Q s^{q-1} \end{aligned}$$

for all $s \in [0, \alpha_Q]$. We also define $\gamma_Q = b_Q \alpha_Q^{q-1}$.

Definition 1 applies to a distribution with support on a subset of \mathbb{R} . The next definition is an extension to distributions defined on a subset of $\mathcal{X} \times \mathbb{R}$.

Definition 2 (Quantiles of p -average type q (Steinwart and Christmann, 2011)). *Let $p \in (0, \infty]$, $q \in [1, \infty)$, and Q be a distribution on $\mathcal{X} \times \mathbb{R}$ with marginal distribution Q_X of \mathcal{X} . Assume that $\text{Supp}(Q(\cdot | X = x)) \subset [-1, 1]$ for Q_X -almost all $x \in X$. Then Q is said to have an α -quantile of p -average type q if $Q(\cdot | X = x)$ has an α -quantile of type q for Q_X -almost all $x \in \mathcal{X}$, and if the function $\gamma : \mathcal{X} \rightarrow [0, \infty]$ given for Q_X -almost all $x \in \mathcal{X}$ by*

$$\gamma(x) := \gamma_{Q(\cdot | X=x)} := b_{Q(\cdot | X=x)} \alpha_{Q(\cdot | X=x)}^{q-1}$$

(as defined in Definition 1) is such that γ^{-1} admits a finite moment $\|\gamma^{-1}\|_{p, Q_X}$ of order p under Q_X .

Now we are ready to state the final assumption we use for the oracle inequality in the sequential setting.

Assumption B6 (Regularity). *For every $(j, \tau) \in \mathcal{J} \times \mathbb{N}^*$, the P_0 -conditional law of $O_{j, \tau}$ given $F_{\tau-1}$ has an α -quantile of p -average type q . Moreover, the collection of $\|\gamma^{-1}\|_{p, Q_X}$, where Q ranges over the P_0 -conditional laws of $O_{j, \tau}$ given $F_{\tau-1}$, is uniformly bounded by a constant $\Gamma > 0$.*

3.1. Super Learner

Discrete Super Learner. Suppose as in the i.i.d. setting that we have K algorithms $\hat{\psi}_1^\alpha, \dots, \hat{\psi}_K^\alpha$ to learn $\psi_{P_0}^\alpha$. Each $\hat{\psi}_k^\alpha$ is a function mapping any finite sequence $\bar{o}_1, \dots, \bar{o}_t$ to an element of Ψ . Let $P_t := \sum_{\tau=1}^t \text{Dirac}(\bar{o}_\tau)$ be the empirical distribution of the data up to time t . Define the L^α -loss of $\psi \in \Psi$ with respect to a batch of observations \bar{o}_t as:

$$\bar{L}^\alpha(\psi)(\bar{o}_t) := \frac{1}{|\mathcal{J}|} \sum_{j \in \mathcal{J}} L^\alpha(\psi)(o_{j,t}).$$

The oracle risk of an algorithm $\hat{\psi}_k^\alpha$ up to time $t \geq 1$ with respect to a distribution $P \in \mathcal{M}$ is defined as:

$$\tilde{R}_{t,P}^\alpha(\hat{\psi}_k^\alpha) := \frac{1}{t} \sum_{\tau=1}^t \mathbb{E}_P \left[\bar{L}^\alpha(\hat{\psi}_k^\alpha(P_{\tau-1}))(\bar{O}_\tau) \middle| F_{\tau-1} \right],$$

where by convention $\hat{\psi}_k^\alpha(P_{t-1})$ produces an arbitrary constant function (e.g. always zero) for the case $t = 1$. The empirical risk up to time t of $\hat{\psi}_k^\alpha$ is defined as the empirical counterpart of its oracle risk:

$$\begin{aligned} \hat{R}_t^\alpha(\hat{\psi}_k^\alpha) &:= \frac{1}{t} \sum_{\tau=1}^t \bar{L}^\alpha(\hat{\psi}_k^\alpha(P_{\tau-1}))(\bar{O}_\tau) \\ &= \frac{1}{t|\mathcal{J}|} \sum_{\tau=1}^t \sum_{j \in \mathcal{J}} L^\alpha(\hat{\psi}_k^\alpha(P_{\tau-1}))(O_{j,\tau}). \end{aligned} \quad (4)$$

At each $t \geq 1$, the discrete online Super Learner selector is formed by finding the algorithm that minimizes the empirical risk up to time t :

$$\hat{\kappa}_t = \arg \min_{k \in [K]} \hat{R}_t^\alpha(\hat{\psi}_k^\alpha).$$

The algorithm $\hat{\psi}_{\hat{\kappa}_t}^\alpha$ is referred to as the online discrete Super Learner.

Continuous Super Learner. Convex combinations of the candidate algorithms are formed as in the i.i.d. setting (3), yielding new candidate algorithms $\hat{\psi}_\pi^\alpha$ (for any $\pi \in \Pi_n$). The continuous online Super Learner selector is then formed by finding weights that minimize the empirical risk in hindsight:

$$\hat{\pi}_t \in \arg \min_{\pi \in \Pi_n} \hat{R}_t^\alpha(\hat{\psi}_\pi^\alpha).$$

The algorithm $\hat{\psi}_{\hat{\pi}_t}^\alpha$ is referred to as the online continuous Super Learner.

Algorithm	R package	Citations
Distributional Random Forest (DRF)	<code>drf</code>	Michel and Cevid (2021); Čevíď et al. (2022)
Gradient Boosting Machine (GBM)	<code>lightgbm</code>	Shi et al. (2023)
Quantile Generalized Additive Models (QGAM)	<code>qgam</code>	Fasiolo et al. (2020, 2021)
Quantile Random Forest (QRF)	<code>grf</code>	Athey et al. (2019)
Quantile Regression (QReg)	<code>quantreg</code>	Koenker (2005); Koenker et al. (2017)
Quantile Regression Neural Network (QRNN)	<code>qrnn</code>	Cannon (2011, 2018)

Table 1: Example library of candidate algorithms for i.i.d. conditional quantile estimation.

3.2. Oracle Inequalities

In this section we establish oracle inequalities for the discrete online Super Learner, based on results for online Super Learner established in (Ecoto et al., 2021). The main result is presented below, with the proof to be found in the appendix.

Theorem 2. *Assume that the number of candidate algorithms grows at most polynomially in t , that is, $K = O(t^a)$ for some $a > 0$. Also assume that, for all $k \in \llbracket K \rrbracket$, $\hat{\psi}_k^\alpha$ only outputs functions which are uniformly bounded by C_0 . Then, for t large enough,*

$$\mathbb{E}_{P_0} \left[\tilde{R}_{t,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_t}^\alpha) - \tilde{R}_{t,P_0}^\alpha(\psi_{P_0}^\alpha) \right] \leq \mathbb{E}_{P_0} \left[\tilde{R}_{t,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_t}^\alpha) - \tilde{R}_{t,P_0}^\alpha(\psi_{P_0}^\alpha) \right] + O\left(\frac{\log(\log(t))}{t^{1/2}}\right).$$

Note that the number of locations $|\mathcal{J}|$ is hidden in the “large enough” t (and not in the order term). Specifically, the larger is the number of locations $|\mathcal{J}|$, the smaller t needs to be for the inequality to hold.

4. Simulation Studies

In this section we investigate the finite sample performance of the QSL in the i.i.d. and sequential settings. A natural application of quantile estimation is in forming prediction intervals from estimates of a lower and upper quantile, which we investigate in both settings. We used the `s13` R package (Coyle et al., 2021) to implement the i.i.d. QSL algorithm. For the sequential setting, we compare the online quantile Super Learner to two algorithms from the online aggregation of experts literature: Exponentially Weighted Average (EWA; Cesa-Bianchi and Lugosi (2006)) and Bernstein Online Aggregation (BOA; Wintenberger (2017)). Both algorithms are implemented in the `opera` R package (Gaillard et al., 2023), which uses an adaptive procedure to fine tune their learning rates. We added the QSL as an additional method to the `opera` package to facilitate its use and comparison with other methods. Code for the simulations and case studies can be found at <https://github.com/herbps10/QuantileSuperLearner>.

4.1. Independent setting

Simulated datasets consisted of $N_1 \in \{250, 500, 1000\}$ i.i.d. draws (X_i, Y_i) , $i = 1, \dots, N_1$, of a generic variable (X, Y) . The covariates forming $X = (X_1, X_2, X_3, X_4, X_5)$ were drawn independently from the uniform distribution on $[0, 1]$. Conditional on X , the outcome Y was chosen to be a linear combination of smooth and non-smooth functions of the covariates:

$$Y = \sin(2X_1) + |X_2| - 0.5X_1X_3 + \lfloor X_4 \rfloor + \epsilon,$$

where $\epsilon \sim N(0, 0.1)$. An additional $N_2 = 1000$ observations (X_i, Y_i) , $i = N_1 + 1, \dots, N_1 + N_2$, were drawn from the same data generating process and used as a validation set to evaluate the performance of the Super Learner in predictions of unseen data. Overall, 50 learning and testing datasets were generated for each sample size N_1 .

We estimated the α -quantile with $\alpha \in \{0.025, 0.05, 0.1, 0.5, 0.9, 0.95, 0.975\}$ for each simulated dataset using separate QSLs. The candidate algorithms included gradient boosting machines, quantile regression, quantile neural networks, quantile random forests, and quantile generalized additive models (see Table 1 for references). We compared the QSL against the candidate algorithms by calculating their empirical risks with respect to the testing dataset. To do so, we defined the empirical risk of an algorithm $\hat{\psi}^\alpha$ as

$$\text{EmpRisk}(\alpha) := \frac{1}{N_2} \sum_{i=N_1+1}^{N_1+N_2} L^\alpha(\hat{\psi}^\alpha(P_{N_1}))(X_i, Y_i). \quad (5)$$

In addition, 80%, 90%, and 95% prediction intervals were formed using the 10%, 5%, and 2.5% and 90%, 95%, and 97.5% quantile estimates as the lower and upper interval bounds, respectively. The empirical coverage of a $(1 - \beta) \times 100\%$ prediction interval built using algorithms $\hat{\psi}^{\beta/2}$ and $\hat{\psi}^{1-\beta/2}$ was defined as

$$\text{EmpCov}(\beta) := \frac{1}{N_2} \sum_{i=N_1+1}^{N_1+N_2} \mathbb{I} \left[\hat{\psi}^{\beta/2}(P_{N_1})(X_i) \leq Y_i \leq \hat{\psi}^{1-\beta/2}(P_{N_1})(X_i) \right]. \quad (6)$$

Results. The empirical risk (5) results are shown in Table 2. QSL achieved the best (or tied for best) quantile risk for all quantiles and sample sizes. However, the empirical coverage prediction intervals formed using QSL estimates, as shown in Table 3, did not necessarily perform as well compared to the candidate algorithms.

4.2. Sequential setting

For the sequential setting we augmented the data generating process from the i.i.d. simulation study to induce temporal dependence. We simulated Y_t , $t = 1, \dots, T = 2000$ following

$$Y_t = \sin(2X_1) + |X_2| - 0.5X_1X_3 + \lfloor X_4 \rfloor + \epsilon_t,$$

where $(\epsilon_t)_{t \leq T}$ is now drawn from an AR(1) process:

$$\epsilon_1 \sim N(0, \sigma^2/(1 - \rho^2)) \quad \text{and, for } 1 < t \leq T, \quad \epsilon_t \sim N(\rho \cdot \epsilon_{t-1}, \sigma^2),$$

with $\rho \in (0, 1)$ an autoregressive parameter and $\sigma > 0$ a scale parameter.

N_1	Algorithm	EmpRisk(α)						
		$\alpha = 0.025$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.975$
250	GRF	0.063	0.11	0.17	0.33	0.17	0.11	0.069
	GBM	0.06	0.084	0.11	0.16	0.1	0.079	0.058
	QGAM	0.034	0.057	0.098	0.29	0.1	0.058	0.035
	QRNN	0.052	0.083	0.13	0.32	0.13	0.083	0.052
	QReg	0.057	0.1	0.17	0.41	0.18	0.1	0.058
	QSL	0.033	0.053	0.088	0.16	0.09	0.057	0.035
500	GRF	0.058	0.095	0.15	0.27	0.16	0.098	0.061
	GBM	0.05	0.065	0.083	0.12	0.084	0.067	0.051
	QGAM	0.032	0.053	0.092	0.27	0.092	0.053	0.032
	QRNN	0.041	0.072	0.12	0.32	0.12	0.073	0.042
	QReg	0.056	0.1	0.17	0.4	0.17	0.1	0.056
	QSL	0.028	0.045	0.07	0.12	0.074	0.049	0.031
1000	GRF	0.051	0.084	0.13	0.2	0.14	0.088	0.054
	GBM	0.038	0.051	0.064	0.092	0.063	0.051	0.04
	QGAM	0.029	0.05	0.088	0.27	0.088	0.05	0.029
	QRNN	0.039	0.068	0.12	0.31	0.12	0.07	0.038
	QReg	0.055	0.099	0.17	0.4	0.17	0.099	0.054
	QSL	0.024	0.038	0.057	0.092	0.058	0.041	0.027

Table 2: Results for the i.i.d. simulation study. The QSL and candidate algorithms were trained on a learning dataset of N_1 observations and evaluated (5) on a testing dataset of $N_2 = 1000$ observations.

N_1	Algorithm	EmpCov(β)		
		$(1 - \beta) = 0.8$	$(1 - \beta) = 0.9$	$(1 - \beta) = 0.95$
250	GRF	88.2%	94.9%	97.2%
	GBM	56.1%	70.7%	81.7%
	QGAM	87.9%	96.9%	99.3%
	QRNN	75.1%	83.7%	87.2%
	QReg	78.8%	88.4%	93.2%
	QSL	82.7%	94.8%	98.3%
500	GRF	89.3%	95.3%	97.6%
	GBM	53%	67.1%	79.3%
	QGAM	89.8%	97.8%	99.7%
	QRNN	77.6%	86.8%	91.5%
	QReg	79.1%	89.1%	93.8%
	QSL	77.2%	92.8%	98.1%
1000	GRF	92%	96.7%	98.5%
	GBM	53.1%	66.5%	78.6%
	QGAM	89.2%	97.6%	99.7%
	QRNN	79.1%	88.6%	93.6%
	QReg	79.9%	89.5%	94.4%
	QSL	73%	89.7%	97.4%

Table 3: Results for the i.i.d. simulation study. The QSL and candidate algorithms were trained on a learning dataset of N_1 observations and evaluated (6) on a testing dataset of $N_2 = 1000$ observations.

We estimated the same set of α -quantiles as before for each simulated dataset using separate online QSLs. The candidate algorithms included gradient boosting machines, quantile regression, quantile neural networks, quantile random forests, and quantile generalized additive models (see Table 1). To save computational time, each of the candidate algorithms was fit once using a single training set of all observations from $t = 1$ to $t = T/2 = 1000$. This amounts to substituting $P_{T/2}$ for $P_{\tau-1}$ in (4).

The empirical risk and coverage of the candidate algorithms are defined as in (5) and (6), substituting $T/2$ for N_1 and N_2 . The empirical risk and coverage of the online algorithms (QSL, EWA, and BOA) is defined differently as they are updated for each t . For an online algorithm $\widehat{\psi}^\alpha$ the final empirical risk is defined as

$$\text{EmpRisk}(\alpha) := \frac{1}{T/2} \sum_{t=T/2+1}^T L^\alpha(\widehat{\psi}_t^\alpha)(X_t, Y_t), \quad (7)$$

where $\widehat{\psi}_t^\alpha$ is the output of the algorithm using all data before time t . Similarly, the final empirical coverage for $(1 - \beta) \times 100\%$ prediction intervals formed from algorithms $\widehat{\psi}^{\beta/2}$ and $\widehat{\psi}^{1-\beta/2}$ is defined as

$$\text{EmpCov}(\beta) := \frac{1}{T/2} \sum_{t=T/2+1}^T \mathbb{I} \left[\widehat{\psi}_t^{\beta/2}(X_t, Y_t) \leq Y_t \leq \widehat{\psi}_t^{1-\beta/2}(X_t, Y_t) \right], \quad (8)$$

where $\widehat{\psi}_t^{\beta/2}$ and $\widehat{\psi}_t^{1-\beta/2}$ are the output of the algorithms $\widehat{\psi}^{\beta/2}$ and $\widehat{\psi}^{1-\beta/2}$ using data before time t , respectively.

Results. The online empirical risk (7) results are shown in Table 4. The algorithms yielded similar empirical risks, with QSL having slightly lower risks for most quantiles and settings of ρ . The prediction intervals formed using the quantile estimates from each algorithm tended to undercover (see Table 5), especially for QSL.

5. Case Studies

In this section we present two case studies based on solar energy applications, one for the i.i.d. setting and one for the online setting.

5.1. Perovskite energy formation and bandgap prediction

A critical component of photovoltaic (PV) cells is the material used for the light-absorbing semiconductor layer. The use of materials with perovskite crystal structures has been the subject of significant recent research, leading to the energy efficiency of perovskite-based PV cells increasing rapidly from 3.8% in 2009 to 26.08% in 2023 (Kojima et al., 2009; Park et al., 2023). While an immense number of compounds exhibit the perovskite crystal structure, not all are useful for solar applications. As experimentally determining the relevant properties of a perovskite compound is resource-intensive, there is significant interest in developing methods to screen for compounds that are likely to have desirable qualities. One method is to use density functional theory (DFT), a method for estimating the physical properties of

ρ	Algorithm	EmpRisk(α)						
		$\alpha = 0.025$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.975$
0	BOA	0.024	0.038	0.056	0.3	0.055	0.04	0.028
	EWA	0.024	0.038	0.057	0.29	0.056	0.04	0.027
	QSL	0.023	0.036	0.055	0.27	0.054	0.038	0.026
0.5	BOA	0.024	0.038	0.058	0.3	0.058	0.041	0.028
	EWA	0.024	0.038	0.059	0.3	0.059	0.041	0.028
	QSL	0.023	0.036	0.056	0.28	0.056	0.039	0.026
0.9	BOA	0.03	0.047	0.073	0.38	0.072	0.048	0.031
	EWA	0.03	0.046	0.072	0.36	0.072	0.048	0.03
	QSL	0.03	0.045	0.072	0.34	0.071	0.047	0.029
0.99	BOA	0.068	0.084	0.14	0.61	0.14	0.082	0.048
	EWA	0.067	0.078	0.13	0.56	0.13	0.078	0.047
	QSL	0.069	0.086	0.14	0.58	0.14	0.083	0.049

Table 4: Results for the online simulation study in terms of empirical risk (7).

ρ	Algorithm	EmpCov(β)		
		$(1 - \beta) = 0.8$	$(1 - \beta) = 0.9$	$(1 - \beta) = 0.95$
0	BOA	71.8%	88.2%	93.5%
	EWA	71.7%	86.9%	92.1%
	QSL	70.9%	86.4%	91.7%
0.5	BOA	72.1%	88.4%	93.6%
	EWA	71.9%	87%	92.1%
	QSL	70.9%	86.5%	92%
0.9	BOA	73.1%	88.7%	92.9%
	EWA	72.8%	87.1%	91.9%
	QSL	71.7%	86.5%	91.5%
0.99	BOA	73.3%	86.3%	87.3%
	EWA	73.9%	86.5%	87.4%
	QSL	71.2%	83.8%	86%

Table 5: Results for the online simulation study in terms of empirical coverage (8).

a compound (Hohenberg and Kohn, 1964; Kohn and Sham, 1965; Jones, 2015). However, the computational intensity of DFT makes it scale poorly to large numbers of candidate perovskites, motivating research into the use of machine learning techniques to approximate DFT outputs as an initial screening step (Chenebua et al., 2021).

In this case study, we focus on predicting the DFT-output *formation energy* and *energy bandgap* of a perovskite material using data made available by Chenebua et al. (2021). Finding the formation energy of a material is useful as it is related to its stability, and determining the energy bandgap is useful as it has a direct relationship with PV efficiency. Chenebua et al. (2021) applied multiple machine learning algorithms to predict formation energy and energy bandgap using data from 1,453 perovskite materials gathered from the Materials Project database (Jain et al., 2013). Predictors include element-based features, stability features, and crystallographic features; we refer to (Chenebua et al., 2021, Table 2) for a full description. Their methods focus on point predictions of formation energy and bandgap; in this case study, we extend their results by forming both point and interval predictions based on quantile estimation.

Formally, let Y_i , $i = 1, \dots, N$ be the DFT formation energy or energy bandgap and \mathbf{X}_i a set of 56 covariates (two of the original covariates from Chenebua et al. (2021) were removed as they were almost perfectly collinear with other covariates). We assume that the couples (X_i, Y_i) , $i = 1, \dots, N$ are drawn i.i.d. from a probability law. Our goal is to estimate the conditional median and conditional α -quantiles (with $\alpha \in \{0.025, 0.05, 0.1, 0.5, 0.9, 0.95, 0.975\}$) of the outcome conditional on covariates. Four candidate algorithms for the Super Learner ensemble were included in the ensemble: generalized random forests (GRF), directional random forests (DRF), gradient boosting machines (GBM), quantile neural networks (QRNN, with 2 hidden layers), and quantile regression (see Table 1).

Results. The case study results are presented in Table 6. For the formation energy and energy bandgap the QSL achieved the lowest cross-validated quantile risk for all quantiles. In addition, the cross-validated empirical coverage of the QSL prediction intervals were the closest to the nominal level. For the energy bandgap, the QSL had the lowest or tied for the lowest empirical risk for 6 out of the 7 estimated quantiles. The QSL 90% and 95% prediction intervals were the closest to having the optimal empirical coverage, although the 80% prediction interval undercovered relative to other methods.

5.2. Post-processing solar irradiance forecasts

Solar irradiance is one of the principal variables influencing photovoltaic power output (Ahmed et al., 2020). Short-term forecasts of solar irradiance are used to predict solar output, which aids electrical grid integration of solar power (Lorenz et al., 2011). As a case study, we apply online quantile Super Learning to generate point predictions and well-calibrated prediction intervals for global horizontal irradiance (GHI) one day in advance, following the case study and data made available by Wang et al. (2022).

As solar irradiance is mainly determined by local meteorological conditions, an important input for GHI forecasts are the outputs of weather prediction models. Numerical Weather Prediction (NWP) is typically based on highly complex deterministic models that output forecasts for a set of meteorological variables over a grid covering a geographical region (or the entire world). Multiple versions of a model are run with slightly perturbed initial conditions

Algorithm	EmpRisk(α)						
	$\alpha = 0.025$	$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 0.9$	$\alpha = 0.95$	$\alpha = 0.975$
<i>Formation Energy</i>							
DRF	0.026	0.044	0.068	0.13	0.079	0.053	0.035
GBM	0.024	0.037	0.049	0.078	0.055	0.042	0.029
GRF	0.029	0.048	0.07	0.12	0.073	0.051	0.032
QRNN	0.02	0.024	0.035	0.079	0.036	0.026	0.015
QReg	0.017	0.026	0.042	0.094	0.044	0.027	0.017
QSL	0.012	0.021	0.034	0.065	0.035	0.021	0.015
<i>Energy Bandgap</i>							
DRF	0.043	0.086	0.17	0.4	0.2	0.13	0.077
GBM	0.047	0.095	0.17	0.3	0.16	0.1	0.067
GRF	0.043	0.086	0.17	0.38	0.19	0.12	0.072
QRNN	0.051	0.089	0.17	0.39	0.18	0.12	0.082
QReg	0.045	0.091	0.18	0.46	0.2	0.12	0.063
QSL	0.043	0.088	0.16	0.3	0.14	0.089	0.054

Table 6: Empirical risk (5) results from the perovskite case study (see Section 5.1). The lowest empirical risk for each task and quantile is bolded. In the case of a tie, all tied algorithms are bolded.

Algorithm	EmpCov(β)		
	$(1 - \beta) = 0.8$	$(1 - \beta) = 0.9$	$(1 - \beta) = 0.95$
<i>Formation Energy</i>			
DRF	94.2%	97.8%	99%
GBM	57.9%	71.6%	84.4%
GRF	94.6%	97.4%	98.7%
QRNN	70.5%	78.2%	82%
QReg	76.1%	86%	89.5%
QSL	79.9%	90.7%	96.4%
<i>Energy Bandgap</i>			
DRF	96.5%	98.5%	99.4%
GBM	59.9%	74.9%	83.7%
GRF	95.6%	98.4%	99.5%
QRNN	74.7%	82%	83.7%
QReg	76.4%	86.2%	88.4%
QSL	68.8%	90.1%	97.2%

Table 7: Empirical coverage (6) results for $(1 - \beta) \times 100\%$ prediction intervals from the perovskite case study (see Section 5.1). The empirical coverage for each task and β value are in bold. In the case of a tie, all tied algorithms are bolded.

to yield a range of plausible weather trajectories which, taken together, are referred to as a *dynamic ensemble* (Du et al., 2019). While dynamic ensemble predictions cover a range of possible scenarios, they are not necessarily well-calibrated in a probabilistic sense (Schulz et al., 2021). For example, in the solar forecasting context, the dispersion of GHI forecasts across the members of a dynamic ensemble may not accurately reflect the variability in the eventually observed GHI. This mismatch motivates post-processing the dynamic ensemble forecasts with an algorithm that produces well-calibrated density estimates or prediction intervals based on the forecasts. If it is not known a-priori which algorithm will perform best for this task, as is almost always the case, then forming ensembles is warranted. It may also be advantageous to regularly update the parameters of the post-processing algorithm with new data as they become available, which leads naturally to the use of online Super Learning.

For this case study, ground truth observations of GHI are taken from satellite measurements available from the National Solar Radiation Data Base (NSRDB; Sengupta et al. (2018)). As NWP input, we use archived forecasts produced by the European Centre for Medium-Range Weather Forecasts (ECMWF) Ensemble Prediction System. Wang et al. (2022) released a subset of historical forecasts covering much of North America and Europe from 2017-2020 at horizons from 0h-90h in advance. Following their case study, and using the case study dataset they released, we focus on predicting GHI at 7 locations in the continental United States using one-day ahead ECMWF forecasts.

Formally, let Y_t , $t = 1, \dots, T$ be the observed satellite measurement of GHI at 13:00 local time on day t at a single location. Let $X_{t,i}$ be the one-day ahead NWP forecasts of GHI covering the same location, where $i = 1, \dots, 50$ indexes the NWP ensemble members. We use as additional covariates the solar zenith angle Z_t and the one-day lagged GHI observation Y_{t-1} . Our goal is to estimate conditional 10%, 50%, and 90% quantiles of GHI separately for each location using quantile Super Learning. To ensure that each of the candidate learners have enough data to produce reasonable predictions, the online learning procedure is started on January 1, 2020, with data from 2017-2019 used as initial training data. Subsequently, each of the candidate algorithms is re-trained after each data point becomes available. The candidate algorithms are as follows (see Table 1 for the R packages we relied on):

- Quantile regression 1: Quantile regression with covariates $X_{t,i}$, $i = 1, \dots, 50$.
- Quantile regression 2: Quantile regression with covariates $X_{t,i}$, $i = 1, \dots, 50$, Z_t , and Y_{t-1} .
- GBM: Gradient Boosting Machines with covariates $X_{t,i}$, $i = 1, \dots, 50$, Z_t , and Y_{t-1} trained with 500 trees.
- GRF: Generalized Random Forests with covariates $X_{t,i}$, $i = 1, \dots, 50$, Z_t , and Y_{t-1} .
- QGAM: Quantile Generalized Additive Models with covariates $X_{t,i}$, $i = 1, \dots, 50$, Z_t , and Y_{t-1} . Spline smooths were used to estimate the association between Z_t and Y_{t-1} and the outcome.
- QRNN: Quantile recurrent neural networks with covariates $X_{t,i}$, $i = 1, \dots, 50$, Z_t , and Y_{t-1} .

α -quantile	Method	EmpRisk(α) by location						
		BON	DRA	FPK	GWN	PSU	SXF	TBL
0.025	BOA	8.66	5.44	7.27	6.53	6.79	6.28	7.87
	EWA	7.83	5.38	7.82	6.91	7.34	6.77	8.04
	QSL	8.04	5.46	7.83	6.88	6.26	6.21	7.24
0.05	BOA	12.7	7.82	11.7	10.9	10.2	10.9	12.1
	EWA	12.7	7.97	12.2	11.3	11.9	11	12.3
	QSL	12.5	7.6	12.3	10.9	9.95	10.2	12.2
0.1	BOA	18.7	10.3	17.5	17.2	16.8	16.3	18.2
	EWA	19	10.6	17.5	17.2	17.1	16	18.2
	QSL	18.5	9.84	17	16.9	16.3	15.6	18.3
0.5	BOA	31.2	13.5	28.5	32.2	31.9	30.5	28
	EWA	31.4	13.5	28.9	32.2	32.1	30.8	28.7
	QSL	31.3	13.5	28.4	32.1	32.1	30.9	28.4
0.9	BOA	15.2	5.34	12.5	14.3	14.2	13.2	11.7
	EWA	15.1	5.57	12.8	14.3	14.3	13.6	12.4
	QSL	14.8	5.47	12.6	13.9	13.9	12.9	11.2
0.95	BOA	9.11	3.36	7.74	8.3	8.44	8.39	7.39
	EWA	9.36	3.11	8.12	8.14	8.16	8.17	7.01
	QSL	8.9	3.06	7.72	8.19	8.1	8.39	7.29
0.975	BOA	5.4	2.39	4.58	5.79	4.54	5.4	4.29
	EWA	5.12	2.4	4.6	5.54	4.9	4.98	4.43
	QSL	5.98	2.14	4.96	5.11	4.38	5.52	4.25

Table 8: Empirical risk of the quantile Super Learner (QSL), Exponentially Weighted Average (EWA), and Bernstein Online Aggregation (BOA) algorithms applied to point forecasting the α -quantiles of ground horizontal irradiance at seven locations in the continental United States (see Section 5.2). The lowest empirical risks for each location are in bold. In the case of a tie, all tied algorithms are bolded. The locations are BON, Bondville, Illinois; DRA, Desert Rock, Nevada; FPK, Fort Peck, Montana; GWN, Goodwin Creek, Mississippi; PSU, Pennsylvania State University, Pennsylvania; SXF, Sioux Falls, South Dakota; and TBL, Table Mountain, Boulder, Colorado.

The performance of the methods are compared by their empirical risk at time T (4). We also evaluate the empirical coverage of the prediction intervals formed from the 10% and 90% quantile forecasts of each method.

Results. The results of applying the quantile Super Learner to the case study dataset are shown in Table 8. The QSL had the lowest empirical risks in most cases for all but the 50% quantile. For the 50% quantile, the QSL tied or had slightly larger empirical risks than the EWA and BOA algorithms. All of the algorithms yielded prediction intervals that were close to the nominal level, with quantile Super Learner having the best-performing intervals (or tied for best) for five of the seven locations.

One way of understanding the empirical performance of each of the candidate algorithms is to examine how they are weighted in the quantile Super Learner ensemble. Figure 1 shows the weights assigned to each candidate algorithm on the final day ($t = T$) for each of the seven locations. Interestingly, the algorithms were weighted differently depending on the

Location	Method	EmpCov(β)		
		$(1 - \beta) = 0.8$	$(1 - \beta) = 0.9$	$(1 - \beta) = 0.95$
BON	BOA	82.8%	91%	95.6%
	EWA	79.2%	89.6%	94.3%
	QSL	80.1%	88.5%	93.2%
DRA	BOA	85.8%	91.8%	95.1%
	EWA	83.1%	90.7%	95.1%
	QSL	81.4%	87.4%	95.1%
FPK	BOA	84.7%	91.3%	93.7%
	EWA	82.8%	89.9%	92.3%
	QSL	80.9%	88.8%	90.7%
GWN	BOA	84.4%	91.3%	95.4%
	EWA	82.8%	89.1%	94.3%
	QSL	81.7%	87.7%	94.8%
PSU	BOA	83.9%	91.3%	93.7%
	EWA	80.3%	88%	92.9%
	QSL	77.3%	86.3%	91.5%
SXF	BOA	84.2%	91.5%	94.3%
	EWA	80.9%	89.6%	92.3%
	QSL	80.6%	88.5%	91%
TBL	BOA	85.2%	91.8%	95.6%
	EWA	80.1%	90.4%	92.6%
	QSL	80.9%	90.7%	92.6%

Table 9: Empirical coverage of the $(1 - \beta) \times 100\%$ prediction intervals of ground horizontal irradiance formed from quantile estimates based on quantile Super Learner (QSL), Exponentially Weighted Average (EWA), and Bernstein Online Aggregation (BOA) at seven locations in the continental United States (see Section 5.2). The empirical coverage closest to the desired level for each location are in bold. In the case of a tie, all tied algorithms are bolded. The locations are BON, Bondville, Illinois; DRA, Desert Rock, Nevada; FPK, Fort Peck, Montana; GWN, Goodwin Creek, Mississippi; PSU, Pennsylvania State University, Pennsylvania; SXF, Sioux Falls, South Dakota; and TBL, Table Mountain, Boulder, Colorado.




Figure 1: Final weights at time $t = T$ assigned to each of the candidate algorithms by the quantile SuperLearner in the solar irradiance forecasting case study (see Section 5.2).

quantile being estimated. For the 10% and 90% quantiles, for example, gradient boosting machines received generally higher weights as compared to the 50% quantile. In general, no single algorithm dominated across all locations and quantiles, illustrating the utility of ensemble based predictions.

6. Discussion

We have proposed a method for conditional quantile estimation, Quantile Super Learning, that combines predictions from multiple candidate algorithms based on their performance measured with respect to a cross-validated empirical risk of the quantile loss function. The approach is theoretically grounded by excess risk bounds that hold with mild assumptions on the data generating distributions in both i.i.d. and sequential data scenarios.

Empirically, in simulation studies the QSL consistently achieved the lowest empirical quantile risk compared to the candidate algorithms in an i.i.d. setup, showing the possible benefit of using ensemble methods. In the sequential setting, we found that online QSL outperformed Exponentially Weighted Average and Bernstein Online Aggregation algorithms in some settings. In the solar irradiance case study (see Section 5.2), QSL tended to achieve lower empirical risk for all but the 50% quantile. Practically QSL is also easy to use as it does not require specifying any tuning parameters, as opposed to EWA and BOA. However, the computational cost of QSL is significantly higher, as an optimization problem must be solved at each time step, as opposed to EWA and BOA in which the weights are updated by simple closed-form equations.

One possible use case for the QSL is to form prediction intervals by separately estimating lower and upper quantiles. However, in both the simulations and the case studies, while Super Learner consistently performed as well as or better than the candidate algorithms in terms of quantile risk, the prediction intervals formed via Super Learning did not always have the best empirical coverage. This reflects the fact that minimizing the quantile loss function does not necessarily lead to optimal coverage. Thus, Super Learner based prediction intervals based on minimizing quantile losses do not enjoy any performance guarantees. A natural extension of this work would be to post-process Super Learner prediction intervals using

techniques from conformal inference, which have strong finite sample results. In the i.i.d. case, conformalized quantile regression or the CV+ method could be used (Romano et al., 2019; Barber et al., 2021). Fakoor et al. (2023) provide a comprehensive empirical comparison of post-processing ensemble quantile estimators, and find conformal inference techniques performed well. In the online case, Adaptive Conformal Inference techniques can be used to endow quantile based prediction intervals with finite sample coverage guarantees (Gibbs and Candès, 2021; Gibbs and Candès, 2022; Zaffran et al., 2022; Bhatnagar et al., 2023). For all of these conformalization approaches it is advantageous to have good underlying estimates of the conditional quantile function, suggesting the use of ensemble methods to hedge against model misspecification.

We note that the goal and implementation of online Super Learning is similar to that of online aggregation of experts approaches in the online learning literature (see the comprehensive overview by Cesa-Bianchi and Lugosi (2006)). Indeed, the online Super Learner functions identically to the Follow the Leader algorithm known to the online learning community. For both approaches, a convex combination of predictions of candidate algorithms is found that minimizes the empirical risk of the ensemble in hindsight. What differentiates them is their theoretical contexts and analyses. Follow the Leader is based on an online learning paradigm which makes no assumptions about how the observed data are generated, including even the possibility of data generated adversarially. It is known to fail in such adversarial settings, and is not favored in the online learning community as there are other algorithms that have better worst-case properties (Cesa-Bianchi and Lugosi, 2006, Chapter 3.2). In the case of i.i.d. data, a now standard online-to-batch argument can be used to translate results about the performance of online learning algorithms, such as regret bounds, to a statistical context. However, these arguments do not apply when the data are dependent, as is expected in many time-series settings. In contrast, the typical analysis of Super Learning is based on a statistical point of view in which the data are posited to follow a probability law for which we assume a stationarity condition, implying that the feature of interest of the conditional law can be learned. Crucially, this rules out the adversarial settings for which Follow the Leader is lacking.

Acknowledgements

This research is partially supported by the Agence Nationale de la Recherche as part of the “Investissements d’avenir” program (reference ANR-19-P3IA-0001; PRAIRIE 3IA Institute).

Competing interests

The authors declare no competing interests.

References

Roger Koenker. *Quantile Regression*. Econometric Society Monographs. Cambridge University Press, 2005. doi: 10.1017/CBO9780511754098.

- Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148 – 1178, 2019. doi: 10.1214/18-AOS1709. URL <https://doi.org/10.1214/18-AOS1709>.
- Alex J. Cannon. Quantile regression neural networks: implementation in r and application to precipitation downscaling. *Computers & Geosciences*, 37:1277–1284. doi:10.1016/j.cageo.2010.07.005, 2011.
- Matteo Fasiolo, Simon N. Wood, Margaux Zaffran, Raphaël Nedellec, and Yannig Goude. Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 116(535):1402–1412, 2020. doi: 10.1080/01621459.2020.1725521. URL <https://doi.org/10.1080/01621459.2020.1725521>.
- Mark J. van der Laan, Eric C Polley, and Alan E. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1), 2007. doi: doi:10.2202/1544-6115.1309. URL <https://doi.org/10.2202/1544-6115.1309>.
- Ziyue Wu and David Benkeser. A huber loss-based super learner with applications to health-care expenditures, 2022. URL <https://arxiv.org/abs/2205.06870>.
- Ingo Steinwart and Andreas Christmann. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1):211 – 225, 2011. doi: 10.3150/10-BEJ267. URL <https://doi.org/10.3150/10-BEJ267>.
- David H. Wolpert. Stacked generalization. *Neural Networks*, 5(2):241–259, 1992. ISSN 0893-6080. doi: [https://doi.org/10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1). URL <https://www.sciencedirect.com/science/article/pii/S0893608005800231>.
- Leo Breiman. Stacked regressions. *Machine Learning*, 24(1):49–64, Jul 1996. ISSN 1573-0565. doi: 10.1007/BF00117832. URL <https://doi.org/10.1007/BF00117832>.
- Mark J. van der Laan and Sandrine Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical Report Working Paper 130, U.C. Berkeley Division of Biostatistics Working Paper Series, 2003. URL <https://biostats.bepress.com/ucbbiostat/paper130>.
- Mark J. van der Laan, Sandrine Dudoit, and Aad W. van der Vaart. The cross-validated adaptive epsilon-net estimator. *Statistics & Decisions*, 24(3):373–395, 2006. doi: doi: 10.1524/stnd.2006.24.3.373. URL <https://doi.org/10.1524/stnd.2006.24.3.373>.
- Aad W. van der Vaart, Sandrine Dudoit, and Mark J. van der Laan. Oracle inequalities for multi-fold cross validation. *Statistics & Decisions*, 24(3):351–371, 2006. doi: doi: 10.1524/stnd.2006.24.3.351. URL <https://doi.org/10.1524/stnd.2006.24.3.351>.
- Erin LeDell, Mark J. van der Laan, and Maya Petersen. Auc-maximizing ensembles through metalearning. *The International Journal of Biostatistics*, 12(1):203–218, 2016. doi: doi: 10.1515/ijb-2015-0035. URL <https://doi.org/10.1515/ijb-2015-0035>.

- David Benkeser, Cheng Ju, Sam Lendle, and Mark van der Laan. Online cross-validation-based ensemble learning. *Statistics in Medicine*, 37(2):249–260, 2018. doi: <https://doi.org/10.1002/sim.7320>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.7320>.
- Geoffrey Ecoto, Aurélien Bibaut, and Antoine Chambaz. One-step ahead sequential super learning from short times series of many slightly dependent data, and anticipating the cost of natural disasters, 2021. URL <https://arxiv.org/abs/2107.13291>.
- Rasool Fakoor, Taesup Kim, Jonas Mueller, Alexander J. Smola, and Ryan J. Tibshirani. Flexible model aggregation for quantile regression. *Journal of Machine Learning Research*, 24(162):1–45, 2023. URL <http://jmlr.org/papers/v24/22-0799.html>.
- Xiaofei Sun, Hongwei Wang, Chao Cai, Mei Yao, and Kangning Wang. Online renewable smooth quantile regression. *Computational Statistics & Data Analysis*, 185:107781, 2023. ISSN 0167-9473. doi: <https://doi.org/10.1016/j.csda.2023.107781>. URL <https://www.sciencedirect.com/science/article/pii/S0167947323000920>.
- Tilmann Gneiting. Quantiles as optimal point forecasts. *International Journal of Forecasting*, 27(2):197–207, 2011. ISSN 0169-2070. doi: <https://doi.org/10.1016/j.ijforecast.2009.12.015>. URL <https://www.sciencedirect.com/science/article/pii/S0169207010000063>.
- Loris Michel and Domagoj Cevic. *drf: Distributional Random Forests*, 2021. URL <https://CRAN.R-project.org/package=drf>. R package version 1.1.0.
- Domagoj Čevic, Loris Michel, Jeffrey Näf, Peter Bühlmann, and Nicolai Meinshausen. Distributional random forests: Heterogeneity adjustment and multivariate distributional regression. *Journal of Machine Learning Research*, 23(333):1–79, 2022. URL <http://jmlr.org/papers/v23/21-0585.html>.
- Yu Shi, Guolin Ke, Damien Soukhavong, James Lamb, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, Tie-Yan Liu, and Nikita Titov. *lightgbm: Light Gradient Boosting Machine*, 2023. URL <https://CRAN.R-project.org/package=lightgbm>. R package version 3.3.5.
- Matteo Fasiolo, Simon N. Wood, Margaux Zaffran, ël Nedellec, and Yannig Goude. qgam: Bayesian nonparametric quantile regression modeling in R. *Journal of Statistical Software*, 100(9):1–31, 2021. doi: 10.18637/jss.v100.i09.
- Roger Koenker, Victor Chernozhukov, Xuming He, and Limin Peng, editors. *Handbook of Quantile Regression*. Chapman and Hall/CRC, 1 edition, 2017. doi: 10.1201/9781315120256.
- Alex J. Cannon. Non-crossing nonlinear regression quantiles by monotone composite quantile regression neural network, with application to rainfall extremes. *Stochastic Environmental Research and Risk Assessment*, 32(11):3207–3225. doi:10.1007/s00477-018-1573-6, 2018.

- Jeremy R Coyle, Nima S Hejazi, Ivana Malenica, Rachael V Phillips, and Oleg Sofrygin. sl3: Modern pipelines for machine learning and Super Learning. <https://github.com/tlverse/sl3>, 2021. URL <https://doi.org/10.5281/zenodo.1342293>. R package version 1.4.2.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006. doi: 10.1017/CBO9780511546921.
- Olivier Wintenberger. Optimal learning with Bernstein online aggregation. *Machine Learning*, 106(1):119–141, Jan 2017. ISSN 1573-0565. doi: 10.1007/s10994-016-5592-6. URL <https://doi.org/10.1007/s10994-016-5592-6>.
- Pierre Gaillard, Yannig Goude, Laurent Plagne, Thibaut Dubois, and Benoit Thieurmél. *opera: Online Prediction by Expert Aggregation*, 2023. URL <http://pierre.gaillard.me/opera.html>. R package version 1.2.1.
- Akihiro Kojima, Kenjiro Teshima, Yasuo Shirai, and Tsutomu Miyasaka. Organometal halide perovskites as visible-light sensitizers for photovoltaic cells. *Journal of the American Chemical Society*, 131(17):6050–6051, 2009. doi: 10.1021/ja809598r. URL <https://doi.org/10.1021/ja809598r>. PMID: 19366264.
- Jaewang Park, Jongbeom Kim, Hyun-Sung Yun, Min Jae Paik, Eunseo Noh, Hyun Jung Mun, Min Gyu Kim, Tae Joo Shin, and Sang Il Seok. Controlled growth of perovskite layers with volatile alkylammonium chlorides. *Nature*, 616(7958):724–730, Apr 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-05825-y. URL <https://doi.org/10.1038/s41586-023-05825-y>.
- P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136:B864–B871, Nov 1964. doi: 10.1103/PhysRev.136.B864. URL <https://link.aps.org/doi/10.1103/PhysRev.136.B864>.
- W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133–A1138, Nov 1965. doi: 10.1103/PhysRev.140.A1133. URL <https://link.aps.org/doi/10.1103/PhysRev.140.A1133>.
- R. O. Jones. Density functional theory: Its origins, rise to prominence, and future. *Rev. Mod. Phys.*, 87:897–923, Aug 2015. doi: 10.1103/RevModPhys.87.897. URL <https://link.aps.org/doi/10.1103/RevModPhys.87.897>.
- Ericsson Tetteh Chenebuah, Michel Nganbe, and Alain Beaudelaire Tchagang. Comparative analysis of machine learning approaches on the prediction of the electronic properties of perovskites: A case study of abx_3 and $a_2bb'x_6$. *Materials Today Communications*, 27:102462, 2021. ISSN 2352-4928. doi: <https://doi.org/10.1016/j.mtcomm.2021.102462>. URL <https://www.sciencedirect.com/science/article/pii/S2352492821004542>.
- Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson. Commentary: The Materials Project: A materials genome approach

- to accelerating materials innovation. *APL Materials*, 1(1), 07 2013. ISSN 2166-532X. doi: 10.1063/1.4812323. URL <https://doi.org/10.1063/1.4812323>. 011002.
- R. Ahmed, V. Sreeram, Y. Mishra, and M.D. Arif. A review and evaluation of the state-of-the-art in pv solar power forecasting: Techniques and optimization. *Renewable and Sustainable Energy Reviews*, 124:109792, 2020. ISSN 1364-0321. doi: <https://doi.org/10.1016/j.rser.2020.109792>. URL <https://www.sciencedirect.com/science/article/pii/S1364032120300885>.
- Elke Lorenz, Thomas Scheidsteger, Johannes Hurka, Detlev Heinemann, and Christian Kurz. Regional pv power prediction for improved grid integration. *Progress in Photovoltaics: Research and Applications*, 19(7):757–771, 2011. doi: <https://doi.org/10.1002/pip.1033>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/pip.1033>.
- Wenting Wang, Dazhi Yang, Tao Hong, and Jan Kleissl. An archived dataset from the ecmwf ensemble prediction system for probabilistic solar power forecasting. *Solar Energy*, 248: 64–75, 2022. ISSN 0038-092X. doi: <https://doi.org/10.1016/j.solener.2022.10.062>. URL <https://www.sciencedirect.com/science/article/pii/S0038092X22008015>.
- Jun Du, Judith Berner, Roberto Buizza, Martin Charron, Peter Houtekamer, Dingchen Hou, Isidora Jankov, Mu Mu, Xuguang Wang, Mozheng Wei, and Huiling Yuan. *Ensemble Methods for Meteorological Predictions*, pages 99–149. Springer Berlin Heidelberg, Berlin, Heidelberg, 2019. ISBN 978-3-642-39925-1. doi: 10.1007/978-3-642-39925-1_13. URL https://doi.org/10.1007/978-3-642-39925-1_13.
- Benedikt Schulz, Mehrez El Ayari, Sebastian Lerch, and Sándor Baran. Post-processing numerical weather prediction ensembles for probabilistic solar irradiance forecasting. *Solar Energy*, 220:1016–1031, 2021. ISSN 0038-092X. doi: <https://doi.org/10.1016/j.solener.2021.03.023>. URL <https://www.sciencedirect.com/science/article/pii/S0038092X21002097>.
- Manajit Sengupta, Yu Xie, Anthony Lopez, Aron Habte, Galen Maclaurin, and James Shelby. The national solar radiation data base (nsrdb). *Renewable and Sustainable Energy Reviews*, 89:51–60, 2018. ISSN 1364-0321. doi: <https://doi.org/10.1016/j.rser.2018.03.003>. URL <https://www.sciencedirect.com/science/article/pii/S136403211830087X>.
- Yaniv Romano, Evan Patterson, and Emmanuel J. Candès. Conformalized quantile regression, 2019.
- Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1):486 – 507, 2021. doi: 10.1214/20-AOS1965. URL <https://doi.org/10.1214/20-AOS1965>.
- Isaac Gibbs and Emmanuel Candès. Adaptive conformal inference under distribution shift. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 1660–1672. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/0d441de75945e5acbc865406fc9a2559-Paper.pdf.

Isaac Gibbs and Emmanuel Candès. Conformal inference for online prediction with arbitrary distribution shifts, 2022.

Margaux Zaffran, Olivier Feron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 25834–25866. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/zaffran22a.html>.

Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning, 2023.

Appendix A. Theoretical Background

In this appendix we present additional theoretical context and proofs of the theorems appearing in the main paper.

Appendix A.1. Independent setting

Our theoretical treatment of the QSL in the i.i.d. setting follows that of Wu and Benkeser (2022), who use results from van der Vaart et al. (2006) and van der Laan et al. (2007). First, we require the definition of a pair of *Bernstein numbers*, on which rests the theoretical analysis.

Definition 3 (van der Vaart et al. (2006)). *Given a measurable $f : \mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}$, and for any $P \in \mathcal{M}$, call $(M(f), v(f))$ a P -pair of Bernstein numbers of f if*

$$M(f)^2 P \left(e^{|f|/M(f)} - 1 - \frac{|f|}{M(f)} \right) \leq \frac{1}{2} v(f).$$

From the following result we see that, for f uniformly bounded, the pair of Bernstein numbers is related to the supremum and variance of f :

Lemma 1 (van der Vaart et al. (2006)). *If f is uniformly bounded then, for any $P \in \mathcal{M}$, $(\|f\|_\infty, 1.5P f^2)$ is a P -pair of Bernstein numbers of f .*

The following result establishes a Bernstein pair for the quantile loss function.

Lemma 2. *Let $\psi \in \Psi$ be uniformly bounded. For any $P \in \mathcal{M}$ let $R_P^\alpha(\psi)$ be the risk of ψ as defined in (1). Under assumptions A1 and A2, the pair $(M(\psi), v(\psi))$ given by*

$$M(\psi) = \max\{\alpha, 1 - \alpha\}(\|\psi\|_\infty + C_0) \quad \text{and} \quad v(\psi) = 1.5 \times M(\psi) \times R_P^\alpha(\psi)$$

is a P -pair of Bernstein numbers of $L^\alpha(\psi)$.

Proof. The proof follows closely that of Wu and Benkeser (2022, Lemma 1); therefore, we only summarize the proof by pointing out the relevant places where it differs. First, see that for all $o = (x, y) \in \mathcal{X} \times [-C_0, C_0]$,

$$\begin{aligned} 0 \leq L^\alpha(\psi)(o) &= \alpha |y - \psi(x)| \mathbb{I}[y > \psi(x)] + (1 - \alpha) |y - \psi(x)| \mathbb{I}[y \leq \psi(x)] \\ &\leq \max\{\alpha, 1 - \alpha\} (\|\psi\|_\infty + C_0) =: M(\psi), \end{aligned}$$

which is the first Bernstein number. Next, we compute $PL^\alpha(\psi)^2$. Following Wu and Benkeser (2022, Lemma 1), we arrive at

$$PL^\alpha(\psi)^2 = M(\psi) \times R_P^\alpha(\psi) =: v(\psi),$$

which is the second Bernstein number. This completes the proof. \square

Next, we state an inequality bounding the difference between the cross-validated risk and the oracle risk.

Theorem 3 (Theorem 2.3, van der Vaart et al. (2006)). *For any $P \in \mathcal{M}$, for any $\psi \in \Psi$ uniformly bounded, let $(M(\psi), v(\psi))$ be a P -pair of Bernstein numbers of the function $L^\alpha(\psi)$. Then for any $\delta > 0$ and $1 \leq p \leq 2$,*

$$\begin{aligned} \mathbb{E}_{P_0} \left[\tilde{R}_{n,P}^\alpha(\hat{\psi}_{\hat{\kappa}_n}^\alpha) \right] &\leq (1 + 2\delta) \mathbb{E}_{P_0} \left[\tilde{R}_{n,P}^\alpha(\hat{\psi}_{\hat{\kappa}_n}^\alpha) \right] \\ &\quad + 16(1 + \delta) \log(1 + K) \\ &\quad \times \mathbb{E}_{B_n} \left[\sup_{\psi \in \Psi} \left(\frac{M(\psi)}{n_1} + \left(\frac{v(\psi)}{n^1 R_P^\alpha(\psi)^{2-p}} \right)^{1/p} \left(\frac{1 + \delta}{\delta} \right)^{2/p-1} \right) \right], \end{aligned}$$

where $n^1 := \sum_{i=1}^n B_n(i)$.

Proof of Theorem 1..

Proof. The proof follows that of Wu and Benkeser (2022, Theorem 2), with the substitution of $\max\{2\alpha C_0, 2(1 - \alpha)C_0\}$ for C . \square

Appendix A.2. Online setting

First, we present the following result that is key to the later analysis.

Theorem 4 (Variance bound for the quantile loss (Steinwart and Christmann, 2011)). *Let $p \in (0, \infty]$, $q \in [0, \infty)$, and*

$$\vartheta = \min \left\{ \frac{2}{q}, \frac{p}{p+1} \right\}.$$

Let Q be a distribution for $O_{j,\tau}$, $(j, \tau) \in \mathcal{J} \times \mathbb{N}^$, that has an α -quantile of p -average type q . For all $\psi \in \Psi$, define $\Delta L^\alpha(\psi) := L^\alpha(\psi) - L^\alpha(\psi_{P_0}^\alpha)$. Then, for all $\psi \in \Psi$, it holds that*

$$Q\Delta L^\alpha(\psi)^2 \leq 2^{2-\vartheta} q^\vartheta \|\gamma^{-1}\|_{p, Q_{X_{j,\tau}}} (Q\Delta L_\alpha(\psi))^\vartheta.$$

In the above statement, $Q_{X_{j,\tau}}$ is the marginal law of $X_{j,\tau}$ under Q . Next, we present several necessary lemmas that follow from Assumptions B1-B6.

Lemma 3. *There exists $b_1 > 0$ such that $\sup_{\psi \in \Psi} \|\Delta L^\alpha(\psi)\|_\infty < b_1$. In addition, there exists $b_2 \in (0, 2b_1]$ such that for all $j \in \mathcal{J}$, $t \geq 1$, and $\psi \in \Psi$, it holds P_0 -almost surely that*

$$|\Delta L^\alpha(\psi)(O_{j,t}) - \mathbb{E}_{P_0}[\Delta L^\alpha(\psi)(O_{j,t}) | F_{t-1}]| \leq b_2.$$

Proof. The existence of b_1 and b_2 follows directly from Assumption B3. \square

Lemma 4. *There exist $\beta \in (0, 1]$ and $\nu > 0$ such that for all $j \in \mathcal{J}$, $t \geq 1$ and $\psi \in \Psi$, it holds P_0 -almost surely that*

$$\mathbb{E}_{P_0} [(\Delta L^\alpha(\psi)(O_{j,t}))^2 | F_{t-1}] \leq \nu (\mathbb{E}_{P_0}[\Delta L^\alpha(\psi)(O_{j,t}) | F_{t-1}])^\beta.$$

Proof. By Assumption B6, the P_0 -conditional laws of $O_{j,t}$ given F_{t-1} have α -quantiles of p -average type q for all $t \geq 1$. Let $\vartheta := \min\left\{\frac{2}{q}, \frac{p}{p+1}\right\}$. By (Theorem 2.8, Steinwart and Christmann, 2011, restated in the Appendix as Theorem 4),

$$\mathbb{E}_{P_0} [(\Delta L^\alpha(\psi)(O_{j,t}))^2 | F_{t-1}] \leq 2^{2-\vartheta} q^\vartheta \|\gamma_0^{-1}\|_p^\vartheta (\mathbb{E}_{P_0}[\Delta L^\alpha(\psi)(O_{j,t}) | F_{t-1}])^\vartheta,$$

where γ_0 is defined as γ in Definition 2 with the choice Q equal to the conditional law of $O_{j,t}$ given F_{t-1} , and $\|\gamma_0^{-1}\|_p$ is the p -norm of γ_0^{-1} with respect to the marginal law of $X_{j,t}$ under Q . In view of the definition of Γ in Assumption B6, setting $\nu = 2^{2-\vartheta} q^\vartheta \Gamma^\vartheta > 0$ and $\beta = \vartheta$ and noting that $0 < \beta < 1$ complete the proof. \square

Lemma 5. *There exists $v_1 > 0$ such that, for all $j \in \mathcal{J}$, $t \geq 1$, and $\psi \in \Psi$, it holds P_0 -almost surely that*

$$\text{Var}[\Delta L^\alpha(\psi)(O_{j,t}) | F_{t-1}] \leq v_1.$$

Proof. The result follows from Lemma 3. \square

Theorem 5 (Oracle Inequality for Online Super Learning (Ecoto et al., 2021)). *Define*

$$v_2 := \frac{3\pi}{2} \left[\left(\frac{15b_2}{|\mathcal{J}|} \right)^2 + \frac{64v_1}{|\mathcal{J}|} \right].$$

For any $\delta \in (0, 1]$, it holds that

$$\begin{aligned} & \mathbb{E}_{P_0} \left[\tilde{R}_{t,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_t}^\alpha) - \tilde{R}_{t,P_0}^\alpha(\psi_{P_0}^\alpha) - (1+2\delta) \left(\tilde{R}_{t,P_0}^\alpha(\hat{\psi}_{\hat{\kappa}_t}^\alpha) - \tilde{R}_{t,P_0}^\alpha(\hat{\psi}_{P_0}^\alpha) \right) \right] \\ & \leq 3 \left(\frac{C_1(\delta)}{t} \log(2KN) \right)^{1/(2-\beta)} + \frac{2C_2(\delta)}{t} \log(2KN), \end{aligned}$$

where $C_1(\delta) := 2^{5-\beta}(1+\delta)^2\gamma/\delta^\beta$, $C_2(\delta) := 8(1+\delta)b_2/3$, and $N \geq 2$ is chosen such that

$$N \geq \frac{\beta}{2-\beta} \frac{\log(t) + \log(C_3)}{\log(2)},$$

with $C_3 := (v_2/\gamma)^{(2-\beta)/\beta}/(2^{5-\beta}\gamma)$.

Proof. It suffices to check each of the 5 conditions of (Ecoto et al., 2021, Corollary 2).

- Assumption 1 is satisfied by our Assumption B4.
- Assumption 2 is satisfied by our Assumption B5.
- Assumption 3 is satisfied according to our Lemma 3.
- Assumption 4 is satisfied according to our Lemma 4.
- Assumption 5 is satisfied because Assumption 3 is satisfied.

□

Proof of Theorem 2.

Proof. To simplify the exposition, let $\text{ER}_{\tilde{\kappa}} := \tilde{R}_{t, P_0}^\alpha(\hat{\psi}_{\tilde{\kappa}_t}^\alpha) - \tilde{R}_{t, P_0}^\alpha(\hat{\psi}_{P_0}^\alpha)$. By Theorem 5,

$$\mathbb{E}_{P_0} \left[\tilde{R}_{t, P_0}^\alpha(\hat{\psi}_{\tilde{\kappa}_t}^\alpha) - \tilde{R}_{t, P_0}^\alpha(\psi_{P_0}^\alpha) \right] \leq \mathbb{E}_{P_0}[\text{ER}_{\tilde{\kappa}}] + 2\delta \mathbb{E}_{P_0}[\text{ER}_{\tilde{\kappa}}] + \text{Rem}(\delta)$$

where

$$\text{Rem}(\delta) := 3 \left(\frac{C_1(\delta)}{t} \log(2KN) \right)^{\frac{1}{2-\beta}} + \frac{2C_2(\delta)}{t} \log(2KN).$$

Choose an integer $N \in [3 \log(t), 4 \log(t)]$, which will necessarily satisfy

$$N \geq \frac{\beta}{2-\beta} \frac{\log(t) + \log(C_3)}{\log(2)}$$

provided that $t \geq \max\{2, C_3\}$. Fix $\delta = t^{-1/2}$. Then

$$C_2(\delta) \leq 16b_2/3 \text{ and } 2^{5-\beta} \gamma t^{\frac{\beta}{2}} \leq C_1(\delta) \leq 2^{7-\beta} \gamma t^{\frac{\beta}{2}}.$$

Therefore, for t large enough,

$$\begin{aligned} \text{Rem}(\delta) &\lesssim C_1(t^{-1/2})^{\frac{1}{2-\beta}} \left[\left(\frac{\log(2KN)}{t} \right)^{\frac{1}{2-\beta}} + \frac{\log(2KN)}{t} \right] \\ &\lesssim C_1(t^{-1/2})^{\frac{1}{2-\beta}} \left[\left(\frac{\log(8K \log(t))}{t} \right)^{\frac{1}{2-\beta}} + \frac{\log(8K \log(t))}{t} \right] \\ &\lesssim \left[C_1(t^{-1/2}) \frac{\log(8K \log(t))}{t} \right]^{\frac{1}{2-\beta}} \\ &\lesssim \left[\frac{\log(8K \log(t))}{t^{1-\frac{\beta}{2}}} \right]^{\frac{1}{2-\beta}} \\ &\lesssim \frac{\log(8K \log(t))}{t^{1/2}}. \end{aligned}$$

Next, note that for large enough t and using the boundedness assumption,

$$2\delta\mathbb{E}[\mathbb{E}R_{\tilde{\kappa}}] = 2t^{-1/2}\mathbb{E}[\mathbb{E}R_{\tilde{\kappa}}] \leq 2t^{-1/2}C_0 \lesssim \frac{\log(8K \log(t))}{t^{1/2}}.$$

Therefore

$$2\delta\mathbb{E}[\mathbb{E}R_{\tilde{\kappa}}] + \text{Rem}(\delta) = O\left(\frac{\log(8K \log(t))}{t^{1/2}}\right),$$

which completes the proof. □