

Is distribution-free inference possible for binary regression?

Rina Foygel Barber

*Department of Statistics
University of Chicago
e-mail: rina@uchicago.edu*

Abstract: For a regression problem with a binary label response, we examine the problem of constructing confidence intervals for the label probability conditional on the features. In a setting where we do not have any information about the underlying distribution, we would ideally like to provide confidence intervals that are distribution-free—that is, valid with no assumptions on the distribution of the data. Our results establish an explicit lower bound on the length of any distribution-free confidence interval, and construct a procedure that can approximately achieve this length. In particular, this lower bound is independent of the sample size and holds for all distributions with no point masses, meaning that it is not possible for any distribution-free procedure to be adaptive with respect to any type of special structure in the distribution.

MSC2020 subject classifications: 62G08, 62G15, 62F35.

Keywords and phrases: Distribution-free, nonparametric inference, binary regression, adaptive inference.

Received May 2020.

1. Introduction

Consider a regression problem where we would like to model the relationship between a feature vector $X \in \mathbb{R}^d$ and a response $Y \in \mathbb{R}$, based on a sample of n data points, $(X_1, Y_1), \dots, (X_n, Y_n) \stackrel{\text{iid}}{\sim} P$. In a high-dimensional setting where d is large, many modern methods are available to build powerful predictive models for Y given X , but relatively little is known about their theoretical properties—for example, if we train a neural network on the n available data points, can we quantify its accuracy on unseen test data, without making strong assumptions on P , the unknown distribution of the data?

If we are willing to assume that the data follows a regression model $\mathbb{E}_P[Y \mid X = x] = f(x)$ where the function f satisfies certain assumptions, then classical statistical results assure that these questions can be answered using more simple regression methods. For example, if $f(x)$ lies in a parametric family (e.g., linear regression) then we can perform inference within this parametric model. In a more general nonparametric setting, if $f(x)$ is assumed to satisfy some smoothness conditions, classical nonparametric methods such as nearest neighbors will also yield guarantees on the accuracy of our estimate of $f(x)$.

However, the reported results will be invalid if the assumptions (the parametric model, or the smoothness conditions) do not hold.

To address this concern, the recent field of *distribution-free prediction* considers the problem of providing valid predictive inference without any assumptions on the data distribution. The aim of distribution-free prediction is formulated as follows: given a training data set $(X_1, Y_1), \dots, (X_n, Y_n) \in \mathbb{R}^d \times \mathbb{R}$, our task is to construct a map \widehat{C}_n , mapping a new data point $x \in \mathbb{R}^d$ to an interval or set $\widehat{C}_n(x) \subseteq \mathbb{R}$, such that

$$\mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ Y_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha \quad \text{for all distributions } P \text{ on } \mathbb{R}^d \times \mathbb{R}. \quad (1)$$

Here the probability is taken with respect to $(X_i, Y_i) \stackrel{\text{iid}}{\sim} P$ for $i = 1, \dots, n+1$ (the training and test data are drawn from the same distribution P). The bound is required to hold uniformly over all distributions P , without constraining to, say, distributions that satisfy some notion of smoothness. For example, the conformal inference methodology (Vovk et al., 2005) provides an elegant framework for distribution-free prediction, and can adapt to the favorable properties of the underlying distribution to achieve asymptotically optimal prediction intervals in certain settings (see, e.g., Lei and Wasserman, 2014; Lei et al., 2018).

Distribution-free prediction has also been studied in the context of a binary response $Y \in \{0, 1\}$, where the output is a set $\widehat{C}_n(X_{n+1}) \subseteq \{0, 1\}$ (or more generally, in a setting with a finite set of possible labels) (Vovk et al., 2005; Lei, 2014; Sadinle et al., 2019). For a binary Y , the goal is to satisfy

$$\mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ Y_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha \quad \text{for all distributions } P \text{ on } \mathbb{R}^d \times \{0, 1\}. \quad (2)$$

For distributions P where the conditional probability $\pi_P(X) = \mathbb{P}_P\{Y = 1 \mid X\}$ is typically close to either 0 or 1, given sufficient data the resulting distribution-free predictive set $\widehat{C}_n(X_{n+1})$ can often be a singleton set, $\{0\}$ or $\{1\}$. If the labels are inherently noisy, however—that is, if $\pi_P(X)$ is typically bounded away from both 0 and 1—then $\{0, 1\}$ will often be the only possible set offering guaranteed predictive coverage, even if we were to have oracle knowledge of the distribution P . In other words, predictive coverage (whether distribution-free or not) is not a meaningful target for binary regression problems with noisy labels; we would like to estimate the label probability $\pi_P(X)$ directly, rather than try to predict the inherently noisy label Y .

1.1. Summary of contributions

In this work, we ask whether the distribution-free framework can be extended beyond the prediction task, in the binary regression setting. We will aim to provide distribution-free inference on the conditional label probability $\pi_P(X) = \mathbb{P}_P\{Y = 1 \mid X\}$. We are particularly interested in scenarios where $\pi_P(X)$ is

typically not close to either 0 or 1 and so meaningful predictive inference would not be possible even if P were known. In this type of setting, is it nonetheless possible to provide a nontrivial confidence interval for $\pi_P(X)$, and to ensure a distribution-free guarantee of coverage?

Specifically, our goal is to investigate the feasibility of constructing an algorithm that satisfies the following condition:

Definition 1. An algorithm \widehat{C}_n provides a $(1 - \alpha)$ -distribution-free confidence interval for binary regression if it holds that

$$\mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ \pi_P(X_{n+1}) \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha$$

for all distributions P on $\mathbb{R}^d \times \{0, 1\}$. (3)

This notion of a valid distribution-free confidence interval was previously studied by Vovk et al. (2005, Section 5.2), under the name “weakly valid probability estimators”.

Definition 1 requires a fairly weak form of coverage—we ask that coverage holds on average over the new feature vector X_{n+1} , rather than requiring $\mathbb{P} \left\{ \pi_P(x) \in \widehat{C}_n(x) \right\} \geq 1 - \alpha$ to hold uniformly over all $x \in \mathbb{R}^d$. Nonetheless, the main results of our work establish that even this weak notion of distribution-free coverage is fundamentally incompatible with the goal of precise inference; with some caveats, the main message of our results is that the property (3) can only be attained by algorithms \widehat{C}_n that return confidence intervals whose length does not vanish with the sample size n . To make this more precise, our main results are the following:

- **Distribution-free confidence leads to distribution-free prediction.** In Theorem 1, we prove that any algorithm \widehat{C}_n satisfying (3) will inevitably also yield a valid prediction interval for Y_{n+1} , for any nonatomic distribution P , i.e., P has no point masses. This result is closely related to Vovk et al. (2005, Proposition 5.1), where it is shown that \widehat{C}_n must include the endpoints 0 and/or 1 with large probability. Intuitively, this implies that, in a noisy setting where $\pi_P(X)$ is not typically close to 0 or 1, any distribution-free confidence interval $\widehat{C}_n(X_{n+1})$ is likely to be quite wide since it needs to reach one or both endpoints.
- **A lower bound on the length of a distribution-free confidence interval.** In Theorem 2, we formalize the above intuition, establishing a lower bound on the expected length of $\widehat{C}_n(X_{n+1})$ with an explicit function of the distribution of $\pi_P(X)$ (again, for any nonatomic P). Importantly, this lower bound is independent of the sample size n . In other words, for any fixed nonatomic distribution P , the length of our distribution-free confidence intervals cannot go to zero even as $n \rightarrow \infty$. This means that distribution-free confidence intervals cannot be adaptive—by requiring coverage to hold for *all* distributions P , we no longer have the possibility of providing precise confidence intervals for *any* distribution P , regardless of whether π_P satisfies “nice” conditions such as smoothness.

- **A matching upper bound.** In Theorem 3 we propose a concrete construction for $\widehat{C}_n(X_{n+1})$ that satisfies the distribution-free coverage property (3). In particular, Corollary 1 proves that the length of our proposed algorithm is asymptotically equal to the lower bound established in Theorem 2, for any distribution P where it is possible to estimate $\pi_P(X)$ consistently as $n \rightarrow \infty$.

1.1.1. Fixed vs. random intervals

In some cases, we may want to allow additional randomness in our construction (formally, we would define \widehat{C}_n as mapping a new feature vector x to a distribution over subsets of \mathbb{R} , and $\widehat{C}_n(x)$ denotes a subset drawn from this distribution). In this case, the probability in statements such as (1), (2), and (3) should be interpreted as being taken with respect to the distribution of the data $(X_i, Y_i) \stackrel{\text{iid}}{\sim} P$ and the additional randomness in the construction of $\widehat{C}_n(X_{n+1})$. From this point on, we will assume that probabilities and expectations are taken on average over any randomness in the construction of the relevant prediction or confidence interval, without further comment. In particular, all results proved in this paper apply to both fixed and random intervals.

1.2. Related work

As mentioned above, the problem of valid distribution-free confidence intervals was previously studied by Vovk et al. (2005, Section 5.2). In addition, this problem is closely related to two lines of work in the recent statistical literature—nonparametric inference (specifically, confidence intervals for nonparametric regression), and distribution-free prediction.

Nonparametric confidence intervals Suppose the response variable Y follows a model $Y = f(X) + \text{noise}$, where $f(X) = \mathbb{E}[Y | X]$ and where the noise distribution is constrained (e.g., subgaussian with some bounded variance). In this setting, we may assume that the true regression function f lies in some constrained class—for example, it may be constrained to be Lipschitz, or to have a Lipschitz gradient (corresponding to a smoothness assumption with exponent $\beta = 1$ or $\beta = 2$, respectively). There is a rich literature on the problems of estimating $f(x)$, and providing inference (e.g., confidence bands) for $f(x)$. If the smoothness level $\beta > 0$ is known, then the problem is fairly straightforward—for example, a k -nearest neighbors method with $k \sim n^{2\beta/(2\beta+d)}$ yields the optimal estimation error rate $\mathcal{O}(n^{-\beta/(2\beta+d)})$, ignoring log factors (and, correspondingly, confidence intervals of this length) (Low, 1997; Györfi et al., 2006).

However, a key question of interest is that of *adaptivity*—if β is unknown but is assumed to satisfy $\beta \geq \beta_0$ is it possible to construct confidence intervals that are valid at smoothness level β_0 , but if applied to data with smoothness $\beta > \beta_0$ would still achieve the optimal rate at that β ? This question has been studied extensively in the literature—see, e.g., Low (1997); Genovese and Wasserman

(2008); Cai et al. (2014) and the references therein. It turns out that the question of adaptivity is closely tied to how we choose to define coverage—if we require coverage at a given point x_0 , i.e., a confidence interval for $f(x_0)$ at a fixed x_0 , then adaptivity is impossible (Low, 1997), while relaxing to nearly-uniform coverage (coverage of $f(x)$ for “most” points x) allows for adaptivity up to $\beta \leq 2\beta_0$ (Cai et al., 2014); a bootstrap based approach for nearly-uniform coverage is studied also by Hall and Horowitz (2013). Adaptivity in the regime $\beta > 2\beta_0$ can be obtained by excluding certain regions of the function space—for instance, under the assumption that the function f is either β -smooth, or is β_0 -smooth and is sufficiently far from any β -smooth function, it becomes possible to detect the correct smoothness level of f and construct the confidence band accordingly. Results of this type are studied by Carpentier (2015); Szabó et al. (2015); Picard and Tribouley (2000); Hoffmann and Nickl (2011); Giné and Nickl (2010); Bull and Nickl (2013). (An overview of many of these results can be found in Giné and Nickl, 2016, Section 8.3.)

A different relaxation of the coverage condition is coverage on average over a random draw of X , studied by Wahba (1983), which is similar to the coverage condition (3) studied in this work. Genovese and Wasserman (2008) propose a different relaxation, providing confidence intervals guaranteed to cover a “surrogate” of the function f (a smoothed version of the regression function). A different notion of coverage is the “confidence ball”, guaranteeing a bound ℓ_2 error $\int_x (\hat{f}(x) - f(x))^2 dx$ rather than providing pointwise confidence intervals for $f(x)$ at each x (see, e.g., Li, 1989; Cai and Low, 2006).

Distribution-free prediction The field of distribution-free prediction aims to provide prediction intervals that are uniformly valid over all distributions, without assuming some minimum level of smoothness as in the nonparametric inference literature. As mentioned above, the conformal prediction framework (Vovk et al., 2005) provides methodology for this aim. Alternative methods offering distribution-free predictive guarantees include holdout set methods (also known as “split” or “inductive” conformal prediction, see, e.g., Papadopoulos et al., 2002; Vovk et al., 2005; Papadopoulos, 2008; Lei et al., 2018), and the jackknife+ (Barber et al., 2019), a variant of the jackknife (i.e., leave-one-out cross-validation). Lei and Wasserman (2014) establish that distribution-free prediction is possible while (approximately) achieving the minimum possible length prediction intervals for any “nice” (e.g., smooth) distribution P .

The work of Vovk (2012); Lei and Wasserman (2014); Barber et al. (2020) study whether a stronger form of predictive coverage can be attained—namely, distribution-free conditional coverage, aiming for a guarantee that holds pointwise at (almost) every x , i.e., $\mathbb{P}\left\{Y_{n+1} \in \widehat{C}_n(X_{n+1}) \mid X_{n+1} = x\right\} \geq 1 - \alpha$. Distribution-free pointwise coverage is shown to be impossible for any finite-length interval (Vovk, 2012; Lei and Wasserman, 2014). Barber et al. (2020) study a weaker notion of conditional coverage, aiming to ensure $\mathbb{P}\left\{Y_{n+1} \in \widehat{C}_n(X_{n+1}) \mid X_{n+1} \in \mathcal{X}\right\} \geq 1 - \alpha$ for all sufficiently large subsets $\mathcal{X} \in \mathbb{R}^d$, and prove lower bounds on the length of any resulting distribution-free

interval. Some of the techniques in these lower bounds are related to the proof techniques we use in the present work.

In the setting where the response Y is binary or takes finitely many values, as discussed earlier, Vovk et al. (2005); Lei (2014); Sadinle et al. (2019) apply the conformal prediction framework to the problem of distribution-free classification. If the goal is to estimate label probabilities (rather than output a predictive set), an alternative notion of validity in the binary setting is *calibration*, where for an estimate $\hat{\pi}(X)$ of the label probability $\pi_P(X)$, we require $\mathbb{P}\{Y = 1 \mid \hat{\pi}(X)\} = \hat{\pi}(X)$. This framework is studied in the distribution-free setting by Vovk and Petej (2014) via the methodology of Venn predictors, and more recently by Gupta et al. (2020).

2. Main results: lower bounds

In this section, we will prove that a distribution-free confidence interval for binary regression cannot provide precise inference about the parameter $\pi_P(X)$. To do this, we first compare to the problem of predictive inference, and then turn to proving lower bounds on the length of any distribution-free confidence interval.

2.1. Confidence vs. prediction

Our first main result proves that, in the binary regression setting, any algorithm providing distribution-free coverage of $\pi_P(X)$ must necessarily also cover the binary label Y , for every distribution P that is nonatomic (i.e., zero probability at any single point).

Theorem 1. *Let \hat{C}_n be any algorithm that provides a $(1 - \alpha)$ -distribution-free confidence interval for binary regression, as in (3). Then \hat{C}_n also satisfies $(1 - \alpha)$ predictive coverage uniformly over all nonatomic distributions P . That is, \hat{C}_n satisfies*

$$\mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ Y_{n+1} \in \hat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha$$

for all nonatomic distributions P on $\mathbb{R}^d \times \{0, 1\}$.

For example, consider a distribution P with a constant label probability, $\pi_P(x) \equiv 0.5$. Given a large sample size n , we might hope that our algorithm would detect the simple nature of this distribution, and could output a narrow interval, $\hat{C}_n(X_{n+1}) = 0.5 \pm o(1)$. However, Theorem 1 tells us that any distribution-free confidence interval \hat{C}_n must necessarily include both endpoints 0 and 1 with substantial probability. In particular, this example suggests that, unless $\pi_P(X)$ is usually close to 0 or 1, any distribution-free confidence interval \hat{C}_n is unlikely to be precise (i.e., $\hat{C}_n(X_{n+1})$ is unlikely to be a short interval). In Theorem 2 below, we will formalize this intuition by finding a lower bound on the expected length of $\hat{C}_n(X_{n+1})$.

We note that Theorem 1 is closely related to a result of Vovk et al. (2005, Proposition 5.1) (see also Nouretdinov et al. (2001, Theorem 7, Corollary 18) for an earlier related result). Their work establishes that if \widehat{C}_n provides a $(1 - \alpha)$ -distribution-free confidence interval for binary regression (3), then there exists some other algorithm \widetilde{C}_n , also satisfying the property (3), such that

$$\mathbb{P} \left\{ \widehat{C}_n(X_{n+1}) \subseteq (0, 1) \right\} \leq \mathbb{P} \left\{ \widetilde{C}_n(X_{n+1}) = \emptyset \right\}.$$

Clearly this last quantity cannot be larger than α (since this would immediately contradict the coverage property (3)), and so this result, like Theorem 1 above, indicates that $\widehat{C}_n(X_{n+1})$ must often include endpoints 0 and/or 1. While Vovk et al. (2005)'s result appears different from the conclusion of Theorem 1 above, the construction in their proof in fact suffices to prove Theorem 1 as well.

2.1.1. A key lemma

Rather than proving Theorem 1 directly, we will instead generalize to a more powerful result:

Lemma 1. *Let \widehat{C}_n be any algorithm that provides a $(1 - \alpha)$ -distribution-free confidence interval for binary regression (3). Let P be any nonatomic distribution on $(X, Y) \in \mathbb{R}^d \times \{0, 1\}$. Then for data points (X_i, Y_i) drawn i.i.d. from P and any random variable $Z_{n+1} \in [0, 1]$ satisfying*

$$Z_{n+1} \perp\!\!\!\perp \widehat{C}_n(X_{n+1}) \mid X_{n+1} \quad \text{and} \quad \mathbb{E}[Z_{n+1} \mid X_{n+1}] = \pi_P(X_{n+1}) \text{ almost surely,}$$

it holds that

$$\mathbb{P} \left\{ Z_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha.$$

(Proofs for this lemma and all subsequent theoretical results are deferred to the Appendix.)

With this lemma in place, Theorem 1 follows immediately—namely, defining $Z_{n+1} = Y_{n+1}$, we have proved the theorem. The lemma is substantially more general, however, and we will need its full generality in order to prove our lower bounds on the length of $\widehat{C}_n(X_{n+1})$ below.

2.2. A lower bound on length

Next, we will establish bounds on the length of a distribution-free confidence interval for binary regression. We begin with a few definitions. First, for $t \in [0, \frac{1}{2}]$ and $a \in [0, 1]$, we define

$$\ell(t, a) = \begin{cases} 2(1 - a)t, & a \geq \frac{1}{2}, \\ \frac{t}{2a}, & a \geq t \text{ and } 0 < a < \frac{1}{2}, \\ 1 - \frac{a}{2t}, & a < t, \\ 0, & a = t = 0, \end{cases}$$

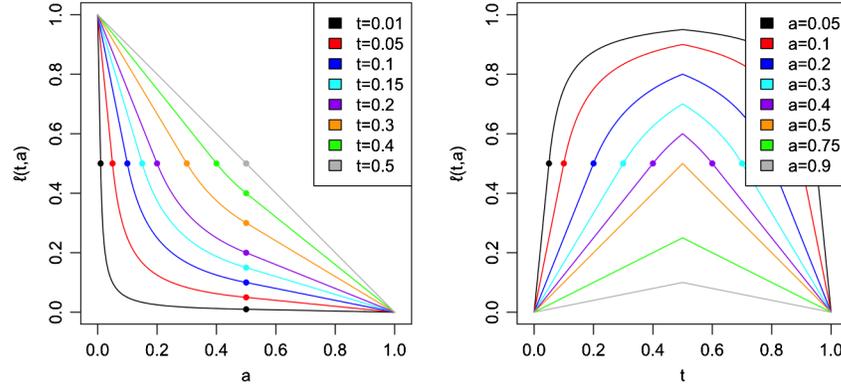


FIG 1. Illustration of the function $\ell(t, a)$. The dots indicate the points where the function switches between the three cases in its definition (i.e., $a = \frac{1}{2}$ and $a = \min\{t, 1 - t\}$).

and for $t \in (\frac{1}{2}, 1]$ let $\ell(t, a) = \ell(1 - t, a)$. The function $\ell(t, a)$ is illustrated in Figure 1. To understand the role of this function in our work, we begin with the following lemma:

Lemma 2. For any $t, a \in [0, 1]$, define

$$\mathcal{F}_{t,a} = \left\{ \begin{array}{l} \text{Measurable functions } f : [0, 1] \rightarrow [0, 1] \text{ satisfying} \\ \mathbb{E}[f(Z)] \geq 1 - a \text{ for any random variable } Z \in [0, 1] \text{ with } \mathbb{E}[Z] = t \end{array} \right\}.$$

Then it holds that

$$\ell(t, a) = \inf_{f \in \mathcal{F}_{t,a}} \left\{ \int_{s=0}^1 f(s) \, ds \right\}.$$

Next, for any distribution Q on $[0, 1]$ and any $\alpha \in [0, 1]$, define

$$L_\alpha(Q) = \inf_{\substack{\text{Measurable fns.} \\ a: [0,1] \rightarrow [0,1]}} \left\{ \mathbb{E}_{T \sim Q} [\ell(T, a(T))] : \mathbb{E}_{T \sim Q} [a(T)] \leq \alpha \right\}.$$

In the following theorem, we will see that this function allows us to explicitly characterize lower and upper bounds for the length of any distribution-free confidence interval.

Theorem 2. Let \widehat{C}_n be any algorithm that provides a $(1 - \alpha)$ -distribution-free confidence interval for binary regression (3). Then for any nonatomic distribution P on $\mathbb{R}^d \times \{0, 1\}$, it holds that

$$\mathbb{E}_{(X_i, Y_i) \stackrel{iid}{\sim} P} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \right] \geq L_\alpha(\Pi_P),$$

where Π_P is the distribution of the random variable $\pi_P(X) \in [0, 1]$.

Here $\text{leb}()$ denotes the Lebesgue measure on \mathbb{R} .

2.2.1. Proof sketch for Theorem 2

First, for intuition, we can consider a simple case where $\Pi_P = \delta_t$, a point mass at some $t \in [0, 1]$ —that is, the label probability is constant, with $\pi_P(x) = t$ for all x . Define a function $f : [0, 1] \rightarrow [0, 1]$ as $f(s) = \mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \{s \in \widehat{C}_n(X_{n+1})\}$. By Lemma 1, for any random variable $Z \in [0, 1]$ with $\mathbb{E}[Z] = t$ drawn independently of the data, it holds that

$$1 - \alpha \leq \mathbb{P} \{Z \in \widehat{C}_n(X_{n+1})\} = \mathbb{E} \left[\mathbb{P} \{Z \in \widehat{C}_n(X_{n+1}) \mid Z\} \right] = \mathbb{E} [f(Z)].$$

Applying Lemma 2, we therefore have

$$\begin{aligned} \ell(t, \alpha) &\leq \int_{s=0}^1 f(s) \, ds = \int_{s=0}^1 \mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \{s \in \widehat{C}_n(X_{n+1})\} \, ds \\ &= \mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\int_{s=0}^1 \mathbb{1} \{s \in \widehat{C}_n(X_{n+1})\} \, ds \right] = \mathbb{E} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \right], \end{aligned}$$

by Fubini’s theorem. We can also verify that $L(\Pi_P) = L(\delta_t) = \ell(t, \alpha)$, completing the proof for this simple case.

Now, in general, $\pi_P(x)$ will not be a constant. Consider any distribution-free confidence interval \widehat{C}_n , and let $a_P(t)$ be the noncoverage rate over test points X_{n+1} conditional on $\pi_P(X_{n+1}) = t$, for a particular distribution P :

$$a_P(t) = \mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ \pi_P(X_{n+1}) \notin \widehat{C}_n(X_{n+1}) \mid \pi_P(X_{n+1}) = t \right\}.$$

We must therefore have $\mathbb{E}_{T \sim \Pi_P} [a_P(T)] \leq \alpha$, in order to achieve at least $1 - \alpha$ coverage. By comparing to the constant-probability case, informally we can see that $\ell(t, a_P(t))$ must be a lower bound on $\mathbb{E} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \mid \pi_P(X_{n+1}) = t \right]$. Therefore, marginalizing over $\pi_P(X_{n+1})$, we see that $\mathbb{E} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \right]$ is lower bounded by $\mathbb{E}_{T \sim \Pi_P} [\ell(T, a_P(T))]$, which is $\geq L_\alpha(\Pi_P)$ by definition. The full details of this proof are deferred to the Appendix.

2.2.2. Interpreting the lower bound

As mentioned in the proof sketch above, we have seen that

$$\text{If } \pi_P(X) = t \text{ almost surely, then } L_\alpha(\Pi_P) = \ell(t, \alpha),$$

giving us an exact expression for the lower bound on length in the case where the label probability is constant.

More generally, we can verify that, for any $t \in [0, \frac{1}{2}]$,

$$\text{If } \pi_P(X) \in [t, 1 - t] \text{ almost surely, then } L_\alpha(\Pi_P) \geq \ell(t, \alpha).$$

(This bound holds because, for any $s \in [t, 1 - t]$, we have $\ell(s, a) \geq \ell(t, a)$ for all a ; examining the definition of L_α leads immediately to this lower bound.) This inequality means that, if $\pi_P(X)$ is bounded away from 0 and 1, then there is a fundamental lower bound on the length of any distribution-free confidence interval regardless of the sample size n , since the lower bound $\ell(t, \alpha) > 0$ does not depend on n . In other words, an infinite sample size does not lead to infinite precision.

2.2.3. Comparison to existing results in predictive inference

Vovk et al. (2005, Section 3.4) study distribution-free prediction in a setting where the label Y takes values in a finite set \mathcal{Y} , with binary labels $Y \in \{0, 1\}$ as a special case. Their results (specifically, see Vovk et al., 2005, Propositions 3.3–3.5) characterize the minimum possible expected cardinality of any distribution-free predictive set in terms of the “predictability” of Y given X —in the special case of a binary label, this translates to studying the distribution Π_P of $\pi_P(X)$. However, the two problems (predictive subsets of $\{0, 1\}$ versus confidence intervals that are subsets of $[0, 1]$) are very different in nature, and their results are not directly related to the map $L_\alpha(\Pi_P)$ derived in our work.

3. Main results: upper bounds

We will next investigate whether the lower bound on confidence interval length, proved in Theorem 2, can in fact be achieved by a distribution-free method. In order to be able to construct confidence intervals based on a finite sample, we will work in a setting where we approximate P via a partition.

We begin by defining some notation. First suppose we are given a predefined partition $\mathbb{R}^d = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_M$. (Later on, for a distribution-free algorithm, we will allow the partition to be chosen as a function of the data.) For each $x \in \mathbb{R}^d$, we define $m(x)$ to be the index of the region containing x , i.e., $\mathcal{X}_{m(x)} \ni x$. For each m , we define

$$p_{P,m} = \mathbb{P}_{P_X} \{X \in \mathcal{X}_m\},$$

the probability of X lying in the m th region, and

$$\pi_{P,m} = \mathbb{E}_{P_X} [\pi_P(X) \mid X \in \mathcal{X}_m] = \mathbb{P}_P \{Y = 1 \mid X \in \mathcal{X}_m\},$$

the average label probability within the m th region.

We will consider confidence intervals that pool data within each region—in particular, we will construct a confidence interval $\widehat{C}_n(X_{n+1})$ that depends on X_{n+1} only through $m(X_{n+1})$. As a result, it is clear that we will only be able to produce a precise confidence interval if $\pi_P(x)$ is approximately constant over $x \in \mathcal{X}_m$, for each m . To capture this notion, we define the “blur” of the partition $\mathcal{X}_{1:M} = \{\mathcal{X}_m\}_{m=1,\dots,M}$ as

$$\Delta_P(\mathcal{X}_{1:M}) = \mathbb{E}_{P_X} [|\pi_P(X) - \pi_{P,m(X)}|].$$

In other words, $\Delta_P(\mathcal{X}_{1:M})$ will be low if the partition is highly informative, separating \mathbb{R}^d into regions where $\pi_P(x)$ is nearly constant. For example, if we have access to a good estimate of the function $\pi_P(x)$, we can partition \mathbb{R}^d by clustering together points with similar estimated probabilities. Of course, the blur $\Delta_P(\mathcal{X}_{1:M})$ cannot in general be small unless we choose M to be large. We will discuss this tradeoff later on.

3.1. An oracle algorithm

In order to motivate our distribution-free construction, we begin with a simpler problem. Suppose that we are given a fixed partition $\mathbb{R}^d = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_M$, and are given oracle knowledge of the probabilities $p_{P,m}$ and $\pi_{P,m}$ defined above. What is the best possible interval length that can be obtained using this oracle knowledge?

As for the lower bound, let us begin by examining the function $\ell(t, a)$. For any $t, a \in [0, 1]$, define a function $f_{t,a} : [0, 1] \rightarrow [0, 1]$ as follows. If $t \in [0, \frac{1}{2}]$, define

$$f_{t,a}(s) = \begin{cases} (1-a) \cdot \mathbb{1}_{s=0}, & \text{if } t = 0, \\ (1-a), & \text{if } t = \frac{1}{2}, \\ 2(1-a) \cdot \max\{1 - \frac{s}{2t}, 0\}, & \text{if } 0 < t < \frac{1}{2} \text{ and } a \geq \frac{1}{2}, \\ \max\{1 - \frac{as}{t}, 0\}, & \text{if } 0 < t < \frac{1}{2} \text{ and } a < \frac{1}{2}, \end{cases} \quad (4)$$

and for $t \in (\frac{1}{2}, 1]$ define $f_{t,a}(s) = f_{1-t,a}(1-s)$. In the proof of Lemma 2, we will establish that $f_{t,a}$ satisfies $\int_{s=0}^1 f_{t,a}(s) ds = \ell(t, a)$, and that $\mathbb{E}[f_{t,a}(Z)] \geq 1-a$ for any Z with $\mathbb{E}[Z] = t$. In other words, the function $f_{t,a}$ attains the infimum in the statement of that lemma.

Next we will leverage this function to construct a confidence interval using the given partition. Fix any $\mathbf{t} = (t_1, \dots, t_M) \in [0, 1]^M$ and any $\mathbf{a} = (a_1, \dots, a_M) \in [0, 1]^M$. Given a test point $x \in \mathbb{R}^d$, we first draw an independent random variable $U \sim \text{Unif}[0, 1]$, and then define

$$C_{\mathbf{t},\mathbf{a}}(x) = \{s \in [0, 1] : f_{t_{m(x)}, a_{m(x)}}(s) \geq U\}. \quad (5)$$

The following lemma examines the length and coverage properties of this construction:

Lemma 3. *Consider a fixed partition $\mathbb{R}^d = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_M$, and a fixed $\mathbf{t} = (t_1, \dots, t_M) \in [0, 1]^M$ and $\mathbf{a} = (a_1, \dots, a_M) \in [0, 1]^M$. Then for any distribution P on $(X, Y) \in \mathbb{R}^d \times \{0, 1\}$,*

$$\mathbb{E}_{P_X} [\text{leb}(C_{\mathbf{t},\mathbf{a}}(X))] = \sum_{m=1}^M p_{P,m} \ell(t_m, a_m).$$

If additionally it holds that

$$\text{For all } m, \text{ either } 0 \leq \pi_{P,m} \leq t_m \leq \frac{1}{2} \text{ or } \frac{1}{2} \leq t_m \leq \pi_{P,m} \leq 1, \quad (6)$$

then

$$\mathbb{P}_{P_X} \{ \pi_P(X) \in C_{\mathbf{t}, \mathbf{a}}(X) \} \geq 1 - \sum_{m=1}^M p_{P,m} a_m.$$

(Note that $C_{\mathbf{t}, \mathbf{a}}(X)$ is a randomized interval, and so the probability and expectation are taken with respect to the data point $X \sim P_X$ and the independent random variable $U \sim \text{Unif}[0, 1]$ used in the construction of $C_{\mathbf{t}, \mathbf{a}}(X)$.)

Next, we define the oracle confidence interval. Suppose we are given oracle knowledge of the $p_{P,m}$'s and $\pi_{P,m}$'s. Define

$$\mathbf{a}_P^* = (a_{P,1}^*, \dots, a_{P,M}^*) \in \underset{a_1, \dots, a_M \in [0,1]}{\operatorname{argmin}} \left\{ \sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_m) : \sum_{m=1}^M p_{P,m} a_m \leq \alpha \right\}. \quad (7)$$

(This is a convex optimization problem, since $\ell(t, a)$ is convex in a .) Given a test point $x \in \mathbb{R}^d$, we define the oracle interval as

$$C_P^*(x) = C_{\boldsymbol{\pi}_P, \mathbf{a}_P^*}(x), \quad (8)$$

where we write $\boldsymbol{\pi}_P = (\pi_{P,1}, \dots, \pi_{P,M})$. To understand this oracle interval, by the results of Lemma 3, we can observe that the constraint $\sum_{m=1}^M p_{P,m} a_m \leq \alpha$ ensures $1 - \alpha$ coverage for the distribution P , while minimizing $\sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_m)$ ensures the lowest possible length under this coverage constraint.

Our next result shows that if the blur $\Delta_P(\mathcal{X}_{1:M})$ of the partition is low, then the expected length of C_P^* is close to the distribution-free lower bound $L_\alpha(\Pi_P)$.

Lemma 4. *The oracle interval C_P^* constructed in (8) satisfies*

$$\mathbb{P}_{P_X} \{ \pi_P(X) \in C_P^* \} \geq 1 - \alpha$$

and

$$\mathbb{E}_{P_X} [\operatorname{leb}(C_P^*(X))] \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}}.$$

Of course, C_P^* is *not* a distribution-free confidence interval—its coverage is guaranteed for a specific distribution P (not uniformly over all P), with the assumption that we have information about the distribution—namely, the $p_{P,m}$'s and $\pi_{P,m}$'s. We next extend this construction into the distribution-free setting by using the training sample to estimate these quantities.

3.2. A distribution-free algorithm

We are now ready to present our distribution-free algorithm. For now, assume again that we are given a fixed partition $\mathbb{R}^d = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_M$ (we will allow for a data-dependent partition later on). After observing a sample $(X_1, Y_1), \dots, (X_n, Y_n) \in \mathbb{R}^d \times \{0, 1\}$, we first estimate the probability $p_{P,m}$ in each region,

$$\hat{p}_m = \frac{\sum_{i=1}^n \mathbb{1}\{X_i \in \mathcal{X}_m\}}{n},$$

and estimate the corresponding label probability $\pi_{P,m}$,

$$\widehat{\pi}_m = \frac{\sum_{i=1}^n \mathbb{1}\{X_i \in \mathcal{X}_m, Y_i = 1\}}{n\widehat{p}_m},$$

or set $\widehat{\pi}_m = \frac{1}{2}$ if $\widehat{p}_m = 0$. In order to ensure distribution-free coverage we will need to work with slightly more conservative estimates. Let

$$\tilde{p}_m = \widehat{p}_m + \sqrt{\widehat{p}_m \cdot \frac{3 \log(4Mn/\alpha)}{n}} + \frac{3 \log(4Mn/\alpha)}{n}, \tag{9}$$

which is chosen to ensure that $p_m \leq \tilde{p}_m$ with high probability, and let

$$\tilde{\pi}_m = \begin{cases} \min \left\{ \frac{1}{2}, \widehat{\pi}_m + \sqrt{\widehat{\pi}_m \cdot \frac{2 \log(4Mn/\alpha)}{n\widehat{p}_m}} + \frac{2 \log(4Mn/\alpha)}{n\widehat{p}_m} \right\}, & \text{if } \widehat{\pi}_m \leq \frac{1}{2}, \\ \max \left\{ \frac{1}{2}, \widehat{\pi}_m - \sqrt{(1 - \widehat{\pi}_m) \cdot \frac{2 \log(4Mn/\alpha)}{n\widehat{p}_m}} - \frac{2 \log(4Mn/\alpha)}{n\widehat{p}_m} \right\}, & \text{if } \widehat{\pi}_m > \frac{1}{2}. \end{cases} \tag{10}$$

This definition of $\tilde{\pi}_m$ is designed to pull the estimate $\widehat{\pi}_m$ closer to $\frac{1}{2}$ (since $\frac{1}{2}$ is the most challenging label probability for coverage, this is therefore a more conservative estimate of $\pi_{P,m}$).

From this point on, we use the same construction (5) as for the ‘‘oracle’’ interval C_P^* above, but with the conservative empirical estimates \tilde{p}_m and $\tilde{\pi}_m$ in place of the unknown true quantities $p_{P,m}$ and $\pi_{P,m}$. Define

$$\tilde{\mathbf{a}} = (\tilde{a}_1, \dots, \tilde{a}_M) \in \operatorname{argmin}_{a_1, \dots, a_M \in [0,1]} \left\{ \sum_{m=1}^M \tilde{p}_m \ell(\tilde{\pi}_m, a_m) : \sum_{m=1}^M \tilde{p}_m a_m \leq \alpha \right\}, \tag{11}$$

and define

$$\widehat{C}_n(X_{n+1}) = C_{\tilde{\pi}, \tilde{\mathbf{a}}}(X_{n+1}), \tag{12}$$

where we write $\tilde{\pi} = (\tilde{\pi}_1, \dots, \tilde{\pi}_M)$.

We now prove that this construction offers a distribution-free confidence interval, and establish an upper bound on its expected length.

Theorem 3. *Let $n \geq 2$. The confidence interval \widehat{C}_n constructed in (12) provides a $(1 - \alpha)$ -distribution-free confidence interval for binary regression (3). Furthermore, for all distributions P on $\mathbb{R}^d \times \{0, 1\}$, the confidence interval \widehat{C}_n satisfies*

$$\mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \right] \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}} + c\sqrt{\frac{M \log n}{\alpha n}},$$

where c is a universal constant.

Comparing to the upper bound calculated in Lemma 4 for the oracle confidence interval C_P^* , the only additional term is $c\sqrt{\frac{M \log n}{\alpha n}}$, which is vanishing with $n \rightarrow \infty$ as long as the partition size M is sufficiently small relative to the sample size n .

3.2.1. Sample splitting for a data-dependent partition

Thus far we have assumed that the partition $\mathbb{R}^d = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_M$ is fixed (and has low blur $\Delta_P(\mathcal{X}_{1:M})$, in order for the upper bound to be meaningful). Of course, in practice, the partition itself would need to be constructed using the data. To do so, we follow a sample splitting strategy. We will use the first half of the training data to estimate the function $\pi_P(x)$ and define the partition accordingly, and the second half of the training data to construct \widehat{C}_n based on this partition. Our construction can be paired with any regression algorithm \mathcal{R} that maps a training data set of size n to an estimate $\widehat{\pi}_n^{\mathcal{R}}(x)$ of the function $\pi_P(x)$ —for example, we might take \mathcal{R} to be logistic regression or k -nearest neighbors. We define the expected error of this regression algorithm as:

$$\Delta_{n,P}(\mathcal{R}) = \mathbb{E} \left[\left| \widehat{\pi}_n^{\mathcal{R}}(X_{n+1}) - \pi_P(X_{n+1}) \right| \right],$$

where the expected value is taken over $(X_1, Y_1), \dots, (X_n, Y_n) \stackrel{\text{iid}}{\sim} P$ and an independent test point $X_{n+1} \sim P_X$. We mention a few special cases:

- In a well-specified parametric model (e.g., $Y|X$ follows a logistic model), we typically have $\Delta_{n,P}(\mathcal{R}) \sim \sqrt{\frac{d}{n}}$ in a low-dimensional ($d < n$) setting, or $\Delta_{n,P}(\mathcal{R}) \sim \sqrt{\frac{k \log d}{n}}$ in a high-dimensional sparse setting, e.g., running ℓ_1 -penalized logistic regression when the true model is k -sparse (Negahban et al., 2012).
- In a nonparametric setting, if $x \mapsto \pi_P(x)$ is assumed to be Lipschitz continuous, a k -nearest neighbors (k -NN) method yields $\Delta_{n,P}(\mathcal{R}) \sim n^{-1/(2+d)}$, or $\Delta_{n,P}(\mathcal{R}) \sim n^{-2/(4+d)}$ if we make the stronger assumption that $x \mapsto \pi_P(x)$ is smooth (see, e.g., Györfi et al., 2006). In the special case where the data is supported on a d_0 -dimensional manifold in \mathbb{R}^d for some $d_0 < d$, then the bound may hold with d_0 in place of d for a faster convergence rate (see, e.g., Jiang, 2019 for finite sample results).

Now we split the sample to construct our distribution-free confidence interval. First, define $\widehat{\pi}_{\lfloor \frac{n}{2} \rfloor}^{\mathcal{R}}$, fitted on data points $i = 1, \dots, \lfloor \frac{n}{2} \rfloor$. Fix $M = \lceil \sqrt{n/\log n} \rceil$, and define

$$\begin{aligned} \widehat{\mathcal{X}}_1^{\mathcal{R}} &= \left\{ x \in \mathbb{R}^d : 0 \leq \widehat{\pi}_{\lfloor \frac{n}{2} \rfloor}^{\mathcal{R}}(x) < \frac{1}{M} \right\}, \dots, \\ \widehat{\mathcal{X}}_M^{\mathcal{R}} &= \left\{ x \in \mathbb{R}^d : \frac{M-1}{M} \leq \widehat{\pi}_{\lfloor \frac{n}{2} \rfloor}^{\mathcal{R}}(x) \leq 1 \right\}. \end{aligned}$$

With this partition in place, we then define \widehat{C}_n exactly as before, except that we restrict our sample to the remaining data points $i = \lfloor \frac{n}{2} \rfloor + 1, \dots, n$. Specifically, let

$$\widehat{p}_m = \frac{\sum_{i=\lfloor \frac{n}{2} \rfloor + 1}^n \mathbb{1} \left\{ X_i \in \widehat{\mathcal{X}}_m^{\mathcal{R}} \right\}}{\lceil \frac{n}{2} \rceil} \quad \text{and} \quad \widehat{\pi}_m = \frac{\sum_{i=\lfloor \frac{n}{2} \rfloor + 1}^n \mathbb{1} \left\{ X_i \in \widehat{\mathcal{X}}_m^{\mathcal{R}}, Y_i = 1 \right\}}{\lceil \frac{n}{2} \rceil \widehat{p}_m},$$

or set $\hat{\pi}_m = \frac{1}{2}$ if $\hat{p}_m = 0$. Define the \tilde{p}_m 's and $\tilde{\pi}_m$'s exactly as in (9) and (10) except with $\lceil \frac{n}{2} \rceil$ in place of n . We then define $\tilde{\mathbf{a}}$ as in (11) before and set

$$\widehat{C}_n^{\mathcal{R}}(X_{n+1}) = C_{\tilde{\pi}, \tilde{\mathbf{a}}}(X_{n+1}). \tag{13}$$

Corollary 1. *Let $n \geq 3$. The interval $\widehat{C}_n^{\mathcal{R}}$ constructed via data splitting with regression algorithm \mathcal{R} (13) provides a $(1-\alpha)$ -distribution-free confidence interval for binary regression (3). Furthermore, for all distributions P on $\mathbb{R}^d \times \{0, 1\}$, $\widehat{C}_n^{\mathcal{R}}$ satisfies*

$$\mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\text{leb}(\widehat{C}_n^{\mathcal{R}}(X_{n+1})) \right] \leq L_\alpha(\Pi_P) + 2\sqrt{\frac{\Delta_{\lfloor \frac{n}{2} \rfloor, P}(\mathcal{R})}{\alpha}} + \frac{c'}{\sqrt{\alpha}} \sqrt[4]{\frac{\log n}{n}},$$

where c' is a universal constant.

This corollary will follow directly from Theorem 3 by observing that, due to the construction of the random partition $\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}$, we have $\mathbb{E} \left[\Delta_P(\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}) \right] \leq 2\Delta_{\lfloor \frac{n}{2} \rfloor, P}(\mathcal{R}) + \frac{1}{M}$.

4. Discussion

The lower bounds established in this paper prove that, in the distribution-free setting, parameter estimation is fundamentally as imprecise as prediction, and confidence intervals for estimating the label probabilities $\pi_P(X) = \mathbb{P}\{Y = 1 \mid X\}$ have a lower bound on their length that does not vanish even with sample size $n \rightarrow \infty$. Unlike the classical literature where these types of results are established for pointwise coverage (i.e., coverage of $\pi_P(x)$ for all x), our new results prove this fundamental lower bound holds even when we require coverage to hold only on average over a new point X drawn from the distribution, and we provide an exact calculation of the minimum possible length. These lower bounds imply that, if we wish to maintain the versatility of the distribution-free setting (i.e., avoiding smoothness assumptions), we can only obtain meaningful confidence intervals by substantially relaxing our notion of coverage. In future work, we hope to examine alternatives—for instance, coverage of a surrogate function approximating $\pi_P(X)$, as in the work of Genovese and Wasserman (2008) for confidence bands in nonparametric regression.

We may also ask whether the results here, proved in the setting of binary regression where $Y \in \{0, 1\}$, may be extended to a more general regression setting—that is, whether it is possible to estimate $\mu_P(X) = \mathbb{E}[Y \mid X]$ under a joint distribution P over $(X, Y) \in \mathbb{R}^d \times \mathbb{R}$. An initial exploration suggests that such an extension would not be straightforward. Specifically, if Y is unbounded, then confidence intervals for $\mu_P(X)$ must necessarily have infinite expected length. To see why, consider a contamination model, where, after drawing $(X, Y) \sim P$, with probability ϵ_n we replace Y with a corrupted value c_n . If we choose $\epsilon_n \ll n^{-1}$, we cannot distinguish between P and this contaminated

model P' on a sample of size n —this means that $\widehat{C}_n(X_{n+1})$ must cover the mean whether the distribution is P or P' . But, by choosing c_n so that $c_n \epsilon_n \rightarrow \infty$, the means $\mu_P(X)$ and $\mu_{P'}(X)$ are arbitrarily far apart, leading to arbitrarily large length of $\widehat{C}_n(X_{n+1})$. Therefore we cannot obtain nontrivial results for unbounded Y . If we instead assume that Y is bounded, on the other hand, then the lower bounds established for the binary case no longer apply (for example, if $Y = 0.5$ almost surely, a distribution-free confidence interval could potentially have vanishing length even if it is constructed only assuming $Y \in [0, 1]$). It is therefore not clear whether there are settings for the general regression problem where we may obtain meaningful upper and lower bounds for distribution-free confidence intervals on the conditional mean $\mu_P(X)$, and we leave these questions for future work.

Appendix A: Additional proofs: lower bounds

A.1. Proof of Lemma 1

The proof of this lemma follows a similar strategy as in Barber et al. (2020, Lemma 3), and is a generalization of the construction used in the proof of Vovk et al. (2005, Proposition 5.1). First we embed the variables in the lemma into a distribution on triples (X, Y, Z) . Writing $\tilde{P}_{X,Z}$ to denote the joint distribution of (X_{n+1}, Z_{n+1}) , we define \tilde{P} as follows:

$$\begin{cases} (X, Z) \sim \tilde{P}_{X,Z}, \\ Y|X, Z \sim \text{Bernoulli}(Z). \end{cases}$$

Note that, marginalizing over Z , the pair (X, Y) follows distribution P . In other words, the joint distribution of

$$(X_1, Y_1), \dots, (X_n, Y_n), X_{n+1}, Z_{n+1}$$

under the model $(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} \tilde{P}$, matches the distribution specified in the lemma. Therefore, it is equivalent to prove the bound

$$\mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} \tilde{P}} \left\{ Z_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha. \quad (14)$$

Next fix any integer $M \geq n + 1$, and let $(X^{(m)}, Z^{(m)}) \stackrel{\text{iid}}{\sim} \tilde{P}_{X,Z}$ for $m = 1, \dots, M$. Let \mathcal{L} specify this sequence of M pairs. Next, fixing \mathcal{L} , we draw $\{(X_i, Y_i, Z_i)\}_{i=1, \dots, n+1}$ as follows:

$$\begin{cases} \text{Sample } m_1, \dots, m_{n+1} \text{ uniformly **without** replacement from } \{1, \dots, M\}, \\ \text{Set } (X_i, Z_i) = (X^{(m_i)}, Z^{(m_i)}) \text{ for each } i = 1, \dots, n + 1, \\ \text{Draw } Y_i \sim \text{Bernoulli}(Z_i) \text{ for each } i = 1, \dots, n + 1. \end{cases} \quad (15)$$

Then clearly, after marginalizing over \mathcal{L} , the triples (X_i, Y_i, Z_i) are drawn i.i.d. from \hat{P} . In other words,

$$\begin{aligned} \mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} \hat{P}} \left\{ Z_{n+1} \in \hat{C}_n(X_{n+1}) \right\} \\ = \mathbb{E}_{\mathcal{L}} \left[\mathbb{P}_{\text{Distrib. (15)}} \left\{ Z_{n+1} \in \hat{C}_n(X_{n+1}) \mid \mathcal{L} \right\} \right]. \end{aligned} \quad (16)$$

Now consider any sequence $\mathcal{L} = \{(X^{(m)}, Z^{(m)})\}_{m=1, \dots, M}$, and let $Q_{\mathcal{L}}$ be the distribution on (X, Y, Z) defined by sampling (X, Z) uniformly at random from \mathcal{L} , then drawing $Y \sim \text{Bernoulli}(Z)$. We can consider drawing $(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} Q_{\mathcal{L}}$ for $i = 1, \dots, n + 1$, which is equivalent to

$$\begin{cases} \text{Sample } m_1, \dots, m_{n+1} \text{ uniformly **with** replacement from } \{1, \dots, M\}, \\ \text{Set } (X_i, Z_i) = (X^{(m_i)}, Z^{(m_i)}) \text{ for each } i = 1, \dots, n + 1, \\ \text{Draw } Y_i \sim \text{Bernoulli}(Z_i) \text{ for each } i = 1, \dots, n + 1. \end{cases} \quad (17)$$

A simple calculation shows that, when m_1, \dots, m_{n+1} are sampled uniformly with replacement from $\{1, \dots, M\}$, the probability of the event $\{m_i = m_j \text{ for any } i \neq j\}$ is bounded by n^2/M . Therefore, for any fixed \mathcal{L} , the total variation distance between the two sampling schemes (15) and (17) is at most n^2/M , and so

$$\begin{aligned} \mathbb{P}_{\text{Distrib. (15)}} \left\{ Z_{n+1} \in \hat{C}_n(X_{n+1}) \mid \mathcal{L} \right\} \\ \geq \mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} Q_{\mathcal{L}}} \left\{ Z_{n+1} \in \hat{C}_n(X_{n+1}) \mid \mathcal{L} \right\} - \frac{n^2}{M}. \end{aligned} \quad (18)$$

Next we calculate $\pi_{Q_{\mathcal{L}}}(x)$, i.e., the label probability under $Q_{\mathcal{L}}$. If $X^{(1)}, \dots, X^{(M)}$ are all distinct for this \mathcal{L} , then by definition of $Q_{\mathcal{L}}$, we can see that $\pi_{Q_{\mathcal{L}}}(X^{(m)}) = Z^{(m)}$ for each $m = 1, \dots, M$. (In particular, we can observe that the probability function $\pi_{Q_{\mathcal{L}}}$ corresponding to this distribution $Q_{\mathcal{L}}$ may in general be highly nonsmooth even if π_P is smooth—we have $|\pi_{Q_{\mathcal{L}}}(X^{(m)}) - \pi_{Q_{\mathcal{L}}}(X^{(m')})| = |Z^{(m)} - Z^{(m')}|$, which may be $\mathcal{O}(1)$ even if $|X^{(m)} - X^{(m')}|$ is arbitrarily small.) Therefore,

$$\begin{aligned} \mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} Q_{\mathcal{L}}} \left\{ Z_{n+1} \in \hat{C}_n(X_{n+1}) \mid \mathcal{L} \right\} \\ = \mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} Q_{\mathcal{L}}} \left\{ \pi_{Q_{\mathcal{L}}}(X_{n+1}) \in \hat{C}_n(X_{n+1}) \mid \mathcal{L} \right\} \geq 1 - \alpha, \end{aligned}$$

where the last step holds since \hat{C}_n is assumed to be an algorithm that provides a $(1 - \alpha)$ -distribution-free confidence interval for any distribution (3), and in particular must provide coverage under $Q_{\mathcal{L}}$. In other words, for any \mathcal{L} , we have proved that

$$\begin{aligned} \mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} Q_{\mathcal{L}}} \left\{ Z_{n+1} \in \hat{C}_n(X_{n+1}) \mid \mathcal{L} \right\} \\ \geq (1 - \alpha) \cdot \mathbb{1} \left\{ X^{(1)}, \dots, X^{(M)} \text{ are distinct} \right\}. \end{aligned}$$

Combining this bound with (16) and (18) establishes that

$$\begin{aligned} \mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} \tilde{P}} \left\{ Z_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \\ \geq (1 - \alpha) \cdot \mathbb{P} \left\{ X^{(1)}, \dots, X^{(M)} \text{ are distinct} \right\} - \frac{n^2}{M}. \end{aligned}$$

Since P is assumed to be nonatomic, the marginal P_X is nonatomic as well, and so the $X^{(m)}$'s are distinct with probability 1. Therefore,

$$\mathbb{P}_{(X_i, Y_i, Z_i) \stackrel{\text{iid}}{\sim} \tilde{P}} \left\{ Z_{n+1} \in \widehat{C}_n(X_{n+1}) \right\} \geq 1 - \alpha - \frac{n^2}{M}.$$

Finally, since M can be taken to be arbitrarily large, this establishes the desired bound (14), and thus completes the proof of the lemma.

A.2. Proof of Lemma 2

We will prove a stronger form of Lemma 2. We define a set of distributions on $[0, 1]$, $\mathcal{Q} = \mathcal{Q}^{(0)} \cup \mathcal{Q}^{(1)}$, where each distribution is a mixture of a point mass and a uniform distribution:

$$\begin{aligned} \mathcal{Q}^{(0)} &= \{p\delta_0 + (1-p)\text{Unif}[0, c] : p, c \in [0, 1]\}, \\ \mathcal{Q}^{(1)} &= \{p\delta_1 + (1-p)\text{Unif}[c, 1] : p, c \in [0, 1]\}. \end{aligned} \tag{19}$$

Lemma 5. For any $t, a \in [0, 1]$, define

$$\mathcal{F}_{t,a}^+ = \left\{ \begin{array}{l} \text{Measurable functions } f : [0, 1] \rightarrow [0, 1] \text{ satisfying} \\ \mathbb{E}[f(Z)] \geq 1 - a \text{ for any random variable } Z \in [0, 1] \text{ such that} \\ \text{either } 0 \leq \mathbb{E}[Z] \leq t \leq \frac{1}{2} \text{ or } \frac{1}{2} \leq t \leq \mathbb{E}[Z] \leq 1 \end{array} \right\}$$

and

$$\mathcal{F}_{t,a}^* = \left\{ \begin{array}{l} \text{Measurable functions } f : [0, 1] \rightarrow [0, 1] \text{ satisfying} \\ \mathbb{E}_Q[f(Z)] \geq 1 - a \text{ for any distribution } Q \in \mathcal{Q} \text{ with } \mathbb{E}_Q[Z] = t \end{array} \right\},$$

and let $\mathcal{F}_{t,a}$ be defined as in the statement of Lemma 2. Then it holds that

$$\ell(t, a) = \inf_{f \in \mathcal{F}_{t,a}^+} \left\{ \int_{s=0}^1 f(s) \, ds \right\} = \inf_{f \in \mathcal{F}_{t,a}} \left\{ \int_{s=0}^1 f(s) \, ds \right\} = \inf_{f \in \mathcal{F}_{t,a}^*} \left\{ \int_{s=0}^1 f(s) \, ds \right\}.$$

Proof of Lemma 5. Since $\mathcal{F}_{t,a}^+ \subseteq \mathcal{F}_{t,a} \subseteq \mathcal{F}_{t,a}^*$, it clearly holds that

$$\inf_{f \in \mathcal{F}_{t,a}^+} \left\{ \int_{s=0}^1 f(s) \, ds \right\} \geq \inf_{f \in \mathcal{F}_{t,a}} \left\{ \int_{s=0}^1 f(s) \, ds \right\} \geq \inf_{f \in \mathcal{F}_{t,a}^*} \left\{ \int_{s=0}^1 f(s) \, ds \right\}.$$

We now need to prove two remaining inequalities to establish the lemma:

$$\inf_{f \in \mathcal{F}_{t,a}^*} \left\{ \int_{s=0}^1 f(s) \, ds \right\} \geq \ell(t, a) \tag{20}$$

and

$$\inf_{f \in \mathcal{F}_{t,a}^+} \left\{ \int_{s=0}^1 f(s) \, ds \right\} \leq \ell(t, a). \tag{21}$$

First we prove (20). Fix any $t, a \in [0, 1]$ and any $f \in \mathcal{F}_{t,a}^*$. We split into cases:

- If $t = 0$, then $\ell(t, a) = 0$, and the bound holds trivially.
- If $0 < t \leq \frac{1}{2} \leq a$, let $Q = \text{Unif}[0, 2t]$. Then $Q \in \mathcal{Q}$ with $\mathbb{E}_Q[Z] = t$, and so we have

$$\int_{s=0}^1 f(s) \, ds \geq 2t \cdot \int_{s=0}^{2t} \frac{1}{2t} \cdot f(s) \, ds = 2t \cdot \mathbb{E}_Q[f(Z)] \geq 2t(1 - a) = \ell(t, a).$$

- If $0 < t \leq a < \frac{1}{2}$, let $Q = (1 - 2a) \cdot \delta_0 + 2a \cdot \text{Unif}[0, \frac{t}{a}]$. Then $Q \in \mathcal{Q}$ with $\mathbb{E}_Q[Z] = t$, and so we have

$$\begin{aligned} \int_{s=0}^1 f(s) \, ds &\geq \frac{t}{2a^2} \cdot 2a \cdot \int_{s=0}^{t/a} \frac{a}{t} \cdot f(s) \, ds \\ &\geq \frac{t}{2a^2} \cdot \left(2a \cdot \int_{s=0}^{t/a} \frac{a}{t} \cdot f(s) \, ds + (1 - 2a) \cdot f(0) - (1 - 2a) \right) \\ &= \frac{t}{2a^2} (\mathbb{E}_Q[f(Z)] - (1 - 2a)) \\ &\geq \frac{t}{2a^2} ((1 - a) - (1 - 2a)) = \frac{t}{2a} = \ell(t, a). \end{aligned}$$

- If $0 \leq a < t \leq \frac{1}{2}$, let $Q = (1 - 2t) \cdot \delta_0 + 2t \cdot \text{Unif}[0, 1]$. Then $Q \in \mathcal{Q}$ with $\mathbb{E}_Q[Z] = t$, and so we have

$$\begin{aligned} \int_{s=0}^1 f(s) \, ds &= \frac{1}{2t} \cdot 2t \cdot \int_{s=0}^1 f(s) \, ds \\ &\geq \frac{1}{2t} \cdot \left(2t \cdot \int_{s=0}^1 f(s) \, ds + (1 - 2t) \cdot f(0) - (1 - 2t) \right) \\ &= \frac{1}{2t} \cdot (\mathbb{E}_Q[f(Z)] - (1 - 2t)) \\ &\geq \frac{1}{2t} \cdot ((1 - a) - (1 - 2t)) = 1 - \frac{a}{2t} = \ell(t, a). \end{aligned}$$

- By symmetry, the analogous calculations hold if $t > \frac{1}{2}$.

Therefore $\int_{s=0}^1 f(s) \, ds \geq \ell(t, a)$ in all cases, which completes the proof of (20).

Next we turn to (21). For any $t, a \in [0, 1]$, define the function $f_{t,a} : [0, 1] \rightarrow [0, 1]$ as in (4). We first verify that

$$\int_{s=0}^1 f_{t,a}(s) \, ds = \ell(t, a). \tag{22}$$

We split into cases:

- If $t = 0$ then $\int_{s=0}^1 f_{t,a}(s) \, ds = 0 = \ell(t, a)$.
- If $t = \frac{1}{2}$ then $\int_{s=0}^1 f_{t,a}(s) \, ds = 1 - a = \ell(t, a)$.
- If $0 < t < \frac{1}{2}$ and $a \geq \frac{1}{2}$, then

$$\int_{s=0}^1 f_{t,a}(s) \, ds = 2(1 - a) \int_{s=0}^{2t} 1 - \frac{s}{2t} \, ds = 2(1 - a)t = \ell(t, a).$$

- If $0 < t \leq a < \frac{1}{2}$, then

$$\int_{s=0}^1 f_{t,a}(s) \, ds = \int_{s=0}^{t/a} 1 - \frac{as}{t} \, ds = \frac{t}{2a} = \ell(t, a).$$

- If $0 \leq a < t < \frac{1}{2}$, then

$$\int_{s=0}^1 f_{t,a}(s) \, ds = \int_{s=0}^1 1 - \frac{as}{t} \, ds = 1 - \frac{a}{2t} = \ell(t, a).$$

- By symmetry, the analogous calculations hold if $t > \frac{1}{2}$.

Therefore we have established that (22) holds in all cases. Next we check that $f_{t,a} \in \mathcal{F}_{t,a}^+$. Let $Z \in [0, 1]$ be any random variable satisfying either $0 \leq \mathbb{E}[Z] \leq t \leq \frac{1}{2}$ or $\frac{1}{2} \leq t \leq \mathbb{E}[Z] \leq 1$. We again split into cases.

- If $t = 0$ then $\mathbb{E}[Z] = 0$ and so $Z = 0$ with probability 1. Then $\mathbb{E}[f_{t,a}(Z)] = f_{0,a}(0) = 1 - a$.
- If $t = \frac{1}{2}$ then $f_{t,a}(Z) = f_{\frac{1}{2},a}(Z) = 1 - a$ almost surely, and so $\mathbb{E}[f_{t,a}(Z)] = 1 - a$.
- If $0 < t < \frac{1}{2}$ and $a \geq \frac{1}{2}$, then $\mathbb{E}[Z] \leq t$ and so

$$\mathbb{E}[f_{t,a}(Z)] = 2(1-a) \cdot \mathbb{E} \left[\max \left\{ 1 - \frac{Z}{2t}, 0 \right\} \right] \geq 2(1-a) \cdot \left(1 - \frac{\mathbb{E}[Z]}{2t} \right) \geq 1 - a.$$

- If $0 < t < \frac{1}{2}$ and $a < \frac{1}{2}$, then $\mathbb{E}[Z] \leq t$ and so

$$\mathbb{E}[f_{t,a}(Z)] = \mathbb{E} \left[\max \left\{ 1 - \frac{aZ}{t}, 0 \right\} \right] \geq 1 - \frac{a\mathbb{E}[Z]}{t} \geq 1 - a.$$

- By symmetry, the analogous calculations hold if $t > \frac{1}{2}$.

Therefore, the bound $\mathbb{E}[f_{t,a}(Z)] \geq 1 - a$ holds in all cases, and so $f_{t,a} \in \mathcal{F}_{t,a}^+$ by definition. Therefore,

$$\inf_{f \in \mathcal{F}_{t,a}^+} \left\{ \int_{s=0}^1 f(s) \, ds \right\} \leq \int_{s=0}^1 f_{t,a}(s) \, ds = \ell(t, a),$$

and so (21) holds. □

A.3. Proof of Theorem 2

Recall the set of distributions \mathcal{Q} defined in (19). Define a function $g : [0, 1] \times [0, 1] \rightarrow [0, 1]$ as

$$g(t, z) = \mathbb{P} \left\{ z \notin \widehat{C}_n(X_{n+1}) \mid \pi_P(X_{n+1}) = t \right\},$$

and define another function $h : [0, 1] \times \mathcal{Q} \rightarrow [0, 1]$ as

$$h(t, Q) = \mathbb{E}_Q [g(t, Z)].$$

For each $t \in [0, 1]$, let $\mathcal{Q}_t = \{Q \in \mathcal{Q} : \mathbb{E}_Q [Z] = t\}$. By Aliprantis and Border (2006, Theorem 18.19), the function

$$t \mapsto a(t) := \sup_{Q \in \mathcal{Q}_t} h(t, Q)$$

is measurable, and furthermore there exists a measurable function $t \mapsto Q_t \in \mathcal{Q}_t$ such that

$$h(t, Q_t) = a(t) = \sup_{Q \in \mathcal{Q}_t} h(t, Q)$$

for all $t \in [0, 1]$, as long as we verify the following conditions:

- \mathcal{Q} is a separable metrizable space. To verify this condition, we will use the total variation distance as our metric on \mathcal{Q} . Define

$$\begin{aligned} \mathcal{Q}_*^{(0)} &= \{p\delta_0 + (1-p)\text{Unif}[0, c] : p, c \in [0, 1] \cap \mathbb{Q}\}, \\ \mathcal{Q}_*^{(1)} &= \{p\delta_1 + (1-p)\text{Unif}[c, 1] : p, c \in [0, 1] \cap \mathbb{Q}\}. \end{aligned}$$

where \mathbb{Q} is the set of rational numbers. Then $\mathcal{Q}_*^{(0)} \cup \mathcal{Q}_*^{(1)}$ is a countable set, and is a dense subset of \mathcal{Q} (under the total variation distance). Therefore, \mathcal{Q} is separable.

- \mathcal{Q}_t is compact for all t . To prove this, first consider the case $t \leq \frac{1}{2}$. If $t = 0$ then $\mathcal{Q}_t \cap \mathcal{Q}^{(0)} = \{\delta_0\}$, and is trivially compact. Now consider $0 < t \leq \frac{1}{2}$. Writing $Q_{p,c} = p\delta_0 + (1-p)\text{Unif}[0, c]$, fix any $Q_{p,c}, Q_{p',c'} \in \mathcal{Q}_t \cap \mathcal{Q}^{(0)}$. We must have $c, c' \geq 2t$ and $p, p' \leq 1 - 2t$ in order to obtain expected value t , and we can calculate that this implies

$$|p - p'| + 2t|c - c'| \leq d_{\text{TV}}(Q_{p,c}, Q_{p',c'}) \leq |p - p'| + (2t)^{-1}|c - c'|.$$

This proves that, on $\mathcal{Q}_t \cap \mathcal{Q}^{(0)}$, the topology induced by total variation distance is the same as the topology induced by the Euclidean distance on $(p, c) \in [0, 1]^2$. Therefore, since $\mathcal{Q}_t \cap \mathcal{Q}^{(0)} = \{Q_{p,c} : (1-p) \cdot \frac{c}{2} = t\}$ corresponds to a closed subset of $(p, c) \in [0, 1]^2$, this set is compact for the case $t \leq \frac{1}{2}$. If instead $t > \frac{1}{2}$, then $\mathcal{Q}_t \cap \mathcal{Q}^{(0)} = \emptyset$ and is trivially compact. An analogous argument shows that $\mathcal{Q}_t \cap \mathcal{Q}^{(1)}$ is compact. Therefore, $\mathcal{Q}_t = (\mathcal{Q}_t \cap \mathcal{Q}^{(0)}) \cup (\mathcal{Q}_t \cap \mathcal{Q}^{(1)})$ is compact.

- $t \mapsto h(t, Q)$ is measurable for all Q , which holds since $t \mapsto g(t, z)$ is measurable by definition of the conditional expectation, and therefore $t \mapsto h(t, Q) = \mathbb{E}_Q [g(t, Z)]$ is measurable as well.
- $Q \mapsto h(t, Q)$ is continuous for all t , which holds since $|h(t, Q) - h(t, Q')| \leq d_{\text{TV}}(Q, Q')$ for all Q, Q' .

Now define a random variable Z_{n+1} drawn from the distribution $Q_{\pi_P(X_{n+1})}$ after conditioning on X_{n+1} . Then by definition, Z_{n+1} satisfies

$$Z_{n+1} \perp\!\!\!\perp \widehat{C}_n(X_{n+1}) \mid X_{n+1} \quad \text{and} \quad \mathbb{E}[Z_{n+1} \mid X_{n+1}] = \pi_P(X_{n+1}) \text{ almost surely,}$$

and by Lemma 1 it therefore holds that $\mathbb{P}\{Z_{n+1} \in \widehat{C}_n(X_{n+1})\} \geq 1 - \alpha$. Therefore,

$$\begin{aligned} \alpha &\geq \mathbb{P}\{Z_{n+1} \notin \widehat{C}_n(X_{n+1})\} \\ &= \mathbb{E}\left[\mathbb{E}\left[\mathbb{P}\{Z_{n+1} \notin \widehat{C}_n(X_{n+1}) \mid Z_{n+1}, \pi_P(X_{n+1})\} \mid \pi_P(X_{n+1})\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}[g(\pi_P(X_{n+1}), Z_{n+1}) \mid \pi_P(X_{n+1})]\right] \\ &= \mathbb{E}[h(\pi_P(X_{n+1}), Q_{\pi_P(X_{n+1})})] = \mathbb{E}[a(\pi_P(X_{n+1}))] = \mathbb{E}_{T \sim \Pi_P}[a(T)]. \end{aligned}$$

Recalling the definition of $L_\alpha(\Pi_P)$, this means that $L_\alpha(\Pi_P) \leq \mathbb{E}_{T \sim \Pi_P}[\ell(T, a(T))]$.

Next, fix any $t \in [0, 1]$. By definition of $a(t)$, for any $Q \in \mathcal{Q}_t$, $\mathbb{E}_Q[g(t, Z)] \leq a(t)$. In the notation of Lemma 5, the function $z \mapsto 1 - g(t, z)$ belongs to $\mathcal{F}_{t, a(t)}^*$, and so

$$\int_{s=0}^1 (1 - g(t, s)) \, ds \geq \ell(t, a(t)).$$

Since this holds for all $t \in [0, 1]$, we therefore have

$$\mathbb{E}_{T \sim \Pi_P} \left[\int_{s=0}^1 (1 - g(T, s)) \, ds \right] \geq \mathbb{E}_{T \sim \Pi_P} [\ell(T, a(T))] \geq L_\alpha(\Pi_P).$$

Finally, applying Fubini's theorem completes the proof:

$$\begin{aligned} \mathbb{E}[\text{leb}(\widehat{C}_n(X_{n+1}))] &= \mathbb{E} \left[\int_{s=0}^1 \mathbb{1}\{s \in \widehat{C}_n(X_{n+1})\} \, ds \right] \\ &= \mathbb{E} \left[\int_{s=0}^1 \mathbb{P}\{s \in \widehat{C}_n(X_{n+1}) \mid \pi_P(X_{n+1})\} \, ds \right] \\ &= \mathbb{E} \left[\int_{s=0}^1 1 - g(\pi_P(X_{n+1}), s) \, ds \right] \\ &= \mathbb{E}_{T \sim \Pi_P} \left[\int_{s=0}^1 1 - g(T, s) \, ds \right] \geq L_\alpha(\Pi_P). \end{aligned}$$

Appendix B: Additional proofs: upper bounds

B.1. Proof of Lemma 3

First we check expected length. By definition of $C_{\mathbf{t},\mathbf{a}}(X_{n+1})$, we have

$$\begin{aligned} \mathbb{E}[\text{leb}(C_{\mathbf{t},\mathbf{a}}(X_{n+1}))] &= \sum_{m=1}^M p_{P,m} \mathbb{E}[\text{leb}(C_{\mathbf{t},\mathbf{a}}(X_{n+1})) \mid X_{n+1} \in \mathcal{X}_m] \\ &= \sum_{m=1}^M p_{P,m} \mathbb{E}[\text{leb}(\{s \in [0, 1] : f_{t_m, a_m}(s) \geq U\})] \\ &= \sum_{m=1}^M p_{P,m} \mathbb{E}\left[\int_{s=0}^1 \mathbb{1}\{f_{t_m, a_m}(s) \geq U\} \, ds\right] \\ &= \sum_{m=1}^M p_{P,m} \int_{s=0}^1 \mathbb{P}\{f_{t_m, a_m}(s) \geq U\} \, ds \quad \text{by Fubini's theorem} \\ &= \sum_{m=1}^M p_{P,m} \int_{s=0}^1 f_{t_m, a_m}(s) \, ds = \sum_{m=1}^M p_{P,m} \ell(t_m, a_m), \end{aligned}$$

where the last step applies the calculation (22) from the proof of Lemma 5.

Next we check coverage under assumption (6). We have

$$\begin{aligned} \mathbb{P}\{\pi_P(X_{n+1}) \in C_{\mathbf{t},\mathbf{a}}(X_{n+1})\} &= \mathbb{P}\{f_{t_m(X_{n+1}), a_m(X_{n+1})}(\pi_P(X_{n+1})) \geq U\} \\ &= \mathbb{E}\left[\mathbb{P}\{f_{t_m(X_{n+1}), a_m(X_{n+1})}(\pi_P(X_{n+1})) \geq U \mid X_{n+1}\}\right] \\ &= \mathbb{E}\left[f_{t_m(X_{n+1}), a_m(X_{n+1})}(\pi_P(X_{n+1}))\right] \\ &= \sum_{m=1}^M p_{P,m} \mathbb{E}\left[f_{t_m, a_m}(\pi_P(X_{n+1})) \mid X_{n+1} \in \mathcal{X}_m\right]. \end{aligned}$$

Next, for each m , in the proof of Lemma 5 we established that $f_{t_m, a_m} \in \mathcal{F}_{t_m, a_m}^+$. By definition of this set, this implies that

$$\mathbb{E}\left[f_{t_m, a_m}(\pi_P(X_{n+1})) \mid X_{n+1} \in \mathcal{X}_m\right] \geq 1 - a_m,$$

since $\mathbb{E}[\pi_P(X_{n+1}) \mid X_{n+1} \in \mathcal{X}_m] = \pi_{P,m}$, and $\pi_{P,m}$ satisfies the assumption (6). Therefore,

$$\mathbb{P}\{\pi_P(X_{n+1}) \in C_{\mathbf{t},\mathbf{a}}(X_{n+1})\} \geq \sum_{m=1}^M p_{P,m} (1 - a_m) = 1 - \sum_{m=1}^M p_{P,m} a_m,$$

as desired.

B.2. Proof of Lemma 4

First, the coverage statement follows immediately from Lemma 3, since $\sum_{m=1}^M p_{P,m} a_{P,m}^* \leq \alpha$ by definition of \mathbf{a}_P^* . Now we consider the expected length. For each x define

$$\delta(x) = |\pi_P(x) - \pi_{P,m(x)}|.$$

Let

$$\epsilon = \mathbb{E}_{P_X} \left[\sqrt{\frac{\delta(X)}{2\alpha}} \right] \leq \sqrt{\frac{\mathbb{E}_{P_X} [\delta(X)]}{2\alpha}} = \sqrt{\frac{\Delta_P(\mathcal{X}_{1:M})}{2\alpha}}. \tag{23}$$

If $\epsilon > 1$ then the result of the lemma holds trivially, since $\text{leb}(C_P^*(X)) \leq \text{leb}([0, 1]) = 1$ always. Therefore we can restrict our attention to the case that $\epsilon \in [0, 1]$.

Now fix any function $a : [0, 1] \rightarrow [0, 1]$ with $\mathbb{E}_{T \sim \Pi_P} [a(T)] \leq \alpha$. Define a vector $\mathbf{a}^\circ \in [0, 1]^M$ with entries

$$a_m^\circ = \min \left\{ 1, \mathbb{E}_{P_X} \left[\sqrt{\frac{\alpha\delta(X)}{2}} + (1 - \epsilon) \cdot a(\pi_P(X)) \mid X \in \mathcal{X}_m \right] \right\}.$$

We can calculate

$$\begin{aligned} \sum_{m=1}^M p_{P,m} a_m^\circ &\leq \sum_{m=1}^M p_{P,m} \mathbb{E}_{P_X} \left[\sqrt{\frac{\alpha\delta(X)}{2}} + (1 - \epsilon) \cdot a(\pi_P(X)) \mid X \in \mathcal{X}_m \right] \\ &= \mathbb{E}_{P_X} \left[\sqrt{\frac{\alpha\delta(X)}{2}} + (1 - \epsilon) \cdot \mathbb{E}_{P_X} [a(\pi_P(X))] \right] \\ &\leq \mathbb{E}_{P_X} \left[\sqrt{\frac{\alpha\delta(X)}{2}} + (1 - \epsilon) \cdot \alpha \right] \text{ by definition of } a \\ &= \alpha \text{ by definition of } \epsilon. \end{aligned}$$

Therefore, \mathbf{a}° is feasible for the optimization problem (7), and so by optimality of \mathbf{a}_P^* , we must have

$$\sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_{P,m}^*) \leq \sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_m^\circ).$$

We next need to bound this right-hand side. For each m , we have either $a_m^\circ = 1$ in which case $\ell(\pi_{P,m}, a_m^\circ) = 0$, or if instead $a_m^\circ < 1$, then we have

$$\begin{aligned} \ell(\pi_{P,m}, a_m^\circ) &= \ell \left(\pi_{P,m}, \mathbb{E}_{P_X} \left[\sqrt{\frac{\alpha\delta(X)}{2}} + (1 - \epsilon) \cdot a(\pi_P(X)) \mid X \in \mathcal{X}_m \right] \right) \\ &\leq \mathbb{E}_{P_X} \left[\ell \left(\pi_{P,m}, \sqrt{\frac{\alpha\delta(X)}{2}} + (1 - \epsilon) \cdot a(\pi_P(X)) \right) \mid X \in \mathcal{X}_m \right] \end{aligned}$$

$$\begin{aligned} &\leq \mathbb{E}_{P_X} \left[\ell(\pi_P(X), (1 - \epsilon) \cdot a(\pi_P(X))) + \frac{|\pi_P(X) - \pi_{P,m}|}{2\sqrt{\frac{\alpha\delta(X)}{2}}} \Bigg| X \in \mathcal{X}_m \right] \\ &= \mathbb{E}_{P_X} \left[\ell(\pi_P(X), (1 - \epsilon) \cdot a(\pi_P(X))) + \sqrt{\frac{\delta(X)}{2\alpha}} \Bigg| X \in \mathcal{X}_m \right], \end{aligned}$$

where the second step applies the fact that $\ell(t, a)$ is convex in a , while the third step applies the fact that $\ell(t, a)$ is $\frac{1}{2a}$ -Lipschitz in t for any fixed a , and is nonincreasing in a for any fixed t . Next, since $\ell(t, a)$ is convex in a and is bounded by 1, we have

$$\begin{aligned} \ell(\pi_P(X), (1 - \epsilon) \cdot a(\pi_P(X))) &= \ell(\pi_P(X), (1 - \epsilon) \cdot a(\pi_P(X)) + \epsilon \cdot 0) \\ &\leq (1 - \epsilon) \cdot \ell(\pi_P(X), a(\pi_P(X))) + \epsilon \cdot \ell(\pi_P(X), 0) \leq \ell(\pi_P(X), a(\pi_P(X))) + \epsilon. \end{aligned}$$

Combining everything, then,

$$\begin{aligned} \ell(\pi_{P,m}, a_m^\circ) &\leq \\ &\mathbb{E}_{P_X} [\ell(\pi_P(X), a(\pi_P(X))) \mid X \in \mathcal{X}_m] + \epsilon + \mathbb{E}_{P_X} \left[\sqrt{\frac{\delta(X)}{2\alpha}} \Bigg| X \in \mathcal{X}_m \right]. \end{aligned}$$

Summing over m , we obtain

$$\begin{aligned} \sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_{P,m}^*) &\leq \sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_m^\circ) \\ &\leq \mathbb{E}_{P_X} [\ell(\pi_P(X), a(\pi_P(X)))] + \epsilon + \mathbb{E}_{P_X} \left[\sqrt{\frac{\delta(X)}{2\alpha}} \right] \\ &\leq \mathbb{E}_{T \sim \Pi_P} [\ell(T, a(T))] + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}}, \end{aligned}$$

where the last step applies (23) and the definition of Π_P . Since we proved this bound for an arbitrary function $a : [0, 1] \rightarrow [0, 1]$ satisfying $\mathbb{E}_{T \sim \Pi_P} [a(T)] \leq \alpha$, we have therefore shown that

$$\sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_{P,m}^*) \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}}.$$

Finally, Lemma 3 implies that $\mathbb{E}_{P_X} [\text{leb}(C_P^*(X))] = \sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_{P,m}^*)$, which completes the proof.

B.3. Proof of Theorem 3

Before proving the theorem, we state two supporting lemmas. The first is a basic property of the function $\ell(t, a)$.

Lemma 6. Fix any $a, t, t' \in [0, 1]$. If it holds that

$$|t - t'| \leq 2\delta^2 + \delta\sqrt{8 \min\{t, 1 - t\}},$$

then

$$\ell(t', a + r\delta) \leq \ell(t, a) + \frac{\delta}{r}$$

for any r such that $a + r\delta \leq 1$.

The second is a simple consequence of the Chernoff bound on the Binomial distribution.

Lemma 7. Let $n \geq 2$. Then under any distribution P , with probability at least $1 - \frac{\alpha}{n}$, the following bounds hold for all $m = 1, \dots, M$:

$$p_{P,m} \leq \tilde{p}_m \left(1 - \frac{1}{n}\right) \text{ and} \\ \tilde{p}_m \leq p_{P,m} + \sqrt{p_{P,m} \cdot \frac{18 \log(4Mn/\alpha)}{n}} + \frac{12 \log(4Mn/\alpha)}{n}, \quad (24)$$

and

$$|\tilde{\pi}_m - \pi_{P,m}| \leq \sqrt{\min\{\pi_{P,m}, 1 - \pi_{P,m}\} \cdot \frac{18 \log(4Mn/\alpha)}{n p_{P,m}}} + \frac{12 \log(4Mn/\alpha)}{n p_{P,m}}, \\ \text{and either } 0 \leq \pi_{P,m} \leq \tilde{\pi}_m \leq \frac{1}{2} \text{ or } \frac{1}{2} \leq \tilde{\pi}_m \leq \pi_{P,m} \leq 1. \quad (25)$$

Now we turn to the proof of the theorem. We first prove the coverage bound under an arbitrary distribution P . Let \mathcal{E}_1 be the event that, for all m , it holds that

$$p_{P,m} \leq \tilde{p}_m \left(1 - \frac{1}{n}\right), \text{ and either } 0 \leq \pi_{P,m} \leq \tilde{\pi}_m \leq \frac{1}{2} \text{ or } \frac{1}{2} \leq \tilde{\pi}_m \leq \pi_{P,m} \leq 1.$$

Lemma 7 verifies that $\mathbb{P}\{\mathcal{E}_1\} \geq 1 - \frac{\alpha}{n}$ for any distribution P . Recall that $\hat{C}_n(X_{n+1}) = C_{\tilde{\pi}, \tilde{\mathbf{a}}}(X_{n+1})$ depends on the training data only through $\tilde{\pi}, \tilde{\mathbf{a}}$, which themselves are functions of the \hat{p}_m 's and $\hat{\pi}_m$'s. By Lemma 3, on the event \mathcal{E}_1 , it holds that

$$\mathbb{P}\left\{\pi_P(X_{n+1}) \notin \hat{C}_n(X_{n+1}) \mid \{(X_i, Y_i)\}_{i=1, \dots, n}\right\} \\ \leq \sum_{m=1}^M p_{P,m} \tilde{a}_m \leq \left(1 - \frac{1}{n}\right) \sum_{m=1}^M \tilde{p}_m \tilde{a}_m \leq \alpha \left(1 - \frac{1}{n}\right),$$

where the last step holds by definition of $\tilde{\mathbf{a}}$. Therefore,

$$\mathbb{P}\left\{\pi_P(X_{n+1}) \notin \hat{C}_n(X_{n+1})\right\} \leq \mathbb{P}\left\{\pi_P(X_{n+1}) \notin \hat{C}_n(X_{n+1}) \mid \mathcal{E}_1\right\} + \frac{\alpha}{n} \leq \alpha,$$

which verifies the distribution-free coverage guarantee.

Next we turn to proving the bound on expected length. Applying Lemma 3, it holds that

$$\mathbb{E} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \mid \{(X_i, Y_i)\}_{i=1, \dots, n} \right] = \sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, \tilde{a}_m).$$

Let \mathcal{E}_2 be the event that, for all m , (24) and (25) both hold. By Lemma 7, $\mathbb{P}\{\mathcal{E}_2\} \geq 1 - \frac{\alpha}{n}$, and therefore,

$$\mathbb{E} \left[\text{leb}(\widehat{C}_n(X_{n+1})) \right] \leq \mathbb{E} \left[\mathbb{1}_{\mathcal{E}_2} \cdot \sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, \tilde{a}_m) \right] + \frac{\alpha}{n}. \tag{26}$$

We therefore now need to bound $\sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, \tilde{a}_m)$ on the event \mathcal{E}_2 .

From this point on, all our calculations will be conditional on $\{(X_i, Y_i)\}_{i=1, \dots, n}$ and we will assume that \mathcal{E}_2 holds. For all m , define

$$\delta_m = \begin{cases} \frac{1}{\sqrt{2}} \left| \sqrt{\tilde{\pi}_m} - \sqrt{\pi_{P,m}} \right|, & \pi_{P,m} \leq \frac{1}{2}, \\ \frac{1}{\sqrt{2}} \left| \sqrt{1 - \tilde{\pi}_m} - \sqrt{1 - \pi_{P,m}} \right|, & \pi_{P,m} > \frac{1}{2}, \end{cases}$$

and let

$$\epsilon = 9 \sqrt{\frac{M \log(4Mn/\alpha)}{\alpha n}}.$$

By (25), we can calculate

$$\sum_{m=1}^M p_{P,m} \delta_m \leq \sum_{m=1}^M p_{P,m} \cdot \sqrt{\frac{6 \log(4Mn/\alpha)}{np_{P,m}}} \leq \sqrt{\frac{6M \log(4Mn/\alpha)}{n}} = \frac{\epsilon \sqrt{6\alpha}}{9}, \tag{27}$$

where the next-to-last step holds since $\sum_{m=1}^M \sqrt{p_{P,m}} \leq \sqrt{M} \sqrt{\sum_{m=1}^M p_{P,m}} = \sqrt{M}$.

Next, we will assume for now that $\epsilon \leq 1$. Let

$$a_m^\circ = \min \{1, (1 - \epsilon) \cdot a_{P,m}^* + \delta_m \sqrt{\alpha}\}$$

where \mathbf{a}_P^* is defined as in (7). Now fix any m . If $a_m^\circ = 1$ then $\ell(\tilde{\pi}_m, a_m^\circ) = 0$. If not, then $a_m^\circ = (1 - \epsilon) \cdot a_{P,m}^* + \delta_m \sqrt{\alpha}$, and we now derive a bound on $\ell(\tilde{\pi}_m, a_m^\circ)$. By definition of δ_m , we have

$$|\tilde{\pi}_m - \pi_{P,m}| \leq \delta_m \cdot \sqrt{8 \min\{\pi_{P,m}, 1 - \pi_{P,m}\}} + 2\delta_m^2.$$

Applying Lemma 6, we have

$$\begin{aligned} \ell(\tilde{\pi}_m, a_m^\circ) &\leq \ell(\pi_{P,m}, a_m^\circ - \delta_m \sqrt{\alpha}) + \frac{\delta_m}{\sqrt{\alpha}} \\ &= \ell(\pi_{P,m}, (1 - \epsilon) a_{P,m}^*) + \frac{\delta_m}{\sqrt{\alpha}} \end{aligned}$$

$$\begin{aligned}
&= \ell(\pi_{P,m}, (1-\epsilon)a_{P,m}^* + \epsilon \cdot 0) + \frac{\delta_m}{\sqrt{\alpha}} \\
&\leq (1-\epsilon) \cdot \ell(\pi_{P,m}, a_{P,m}^*) + \epsilon \cdot \ell(\pi_{P,m}, 0) + \frac{\delta_m}{\sqrt{\alpha}} \\
&\leq \ell(\pi_{P,m}, a_{P,m}^*) + \epsilon + \frac{\delta_m}{\sqrt{\alpha}},
\end{aligned}$$

since $\ell(t, a)$ is convex in a , and bounded by 1. Summing over all m , and applying (27),

$$\begin{aligned}
\sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, a_m^\circ) &\leq \sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_{P,m}^*) + \left(1 + \frac{\sqrt{6}}{9}\right) \epsilon \\
&\leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}} + \left(1 + \frac{\sqrt{6}}{9}\right) \epsilon, \quad (28)
\end{aligned}$$

where the last step applies Lemma 4 along with the calculation $\mathbb{E}_{P_X}[\text{leb}(C_P^*(X))]$ = $\sum_{m=1}^M p_{P,m} \ell(\pi_{P,m}, a_{P,m}^*)$ from the proof of that lemma.

Next we will need to relate $\sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, a_m^\circ)$ to $\sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, \tilde{a}_m)$. Recalling the definition of $\tilde{\mathbf{a}}$ in (11), we see that by optimality of $\tilde{\mathbf{a}}$,

$$\text{If } \sum_{m=1}^M \tilde{p}_m a_m^\circ \leq \alpha, \text{ then } \sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, \tilde{a}_m) \leq \sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, a_m^\circ). \quad (29)$$

We now turn to verifying that $\sum_{m=1}^M \tilde{p}_m a_m^\circ \leq \alpha$, to ensure that \mathbf{a}° is feasible for the optimization problem (11). First, we have

$$\begin{aligned}
\sum_{m=1}^M p_{P,m} a_m^\circ &\leq \sum_{m=1}^M p_{P,m} ((1-\epsilon) \cdot a_{P,m}^* + \delta_m \sqrt{\alpha}) \\
&= (1-\epsilon) \sum_{m=1}^M p_{P,m} a_{P,m}^* + \sqrt{\alpha} \sum_{m=1}^M p_{P,m} \delta_m \\
&\leq (1-\epsilon)\alpha + \sqrt{\alpha} \sum_{m=1}^M p_{P,m} \delta_m \quad \text{by definition of } \mathbf{a}_P^* \quad (7) \\
&\leq (1-\epsilon)\alpha + \sqrt{\alpha} \cdot \frac{\epsilon\sqrt{6}\alpha}{9} \quad \text{by (27)} \\
&= \alpha - \alpha\epsilon \left(1 - \frac{\sqrt{6}}{9}\right).
\end{aligned}$$

Next, applying (24), along with the fact that $a_m^\circ \leq 1$ by construction, we have

$$\sum_{m=1}^M \tilde{p}_m a_m^\circ \leq \sum_{m=1}^M \left(p_{P,m} + \sqrt{p_{P,m} \cdot \frac{18 \log(4Mn/\alpha)}{n}} + \frac{12 \log(4Mn/\alpha)}{n} \right) \cdot a_m^\circ$$

$$\begin{aligned}
 &\leq \sum_{m=1}^M p_{P,m} a_m^\circ + \left(\sum_{m=1}^M \sqrt{p_{P,m} a_m^\circ} \right) \cdot \sqrt{\frac{18 \log(4Mn/\alpha)}{n}} + \frac{12M \log(4Mn/\alpha)}{n} \\
 &\leq \sum_{m=1}^M p_{P,m} a_m^\circ + \sqrt{\sum_{m=1}^M p_{P,m} a_m^\circ} \cdot \sqrt{\frac{18M \log(4Mn/\alpha)}{n}} + \frac{12M \log(4Mn/\alpha)}{n} \\
 &\leq \alpha - \alpha\epsilon \left(1 - \frac{\sqrt{6}}{9} \right) + \sqrt{\alpha} \cdot \sqrt{\frac{18M \log(4Mn/\alpha)}{n}} + \frac{12M \log(4Mn/\alpha)}{n} \\
 &\leq \alpha,
 \end{aligned}$$

where the last step plugs in the definition of ϵ along with the assumption that $\epsilon \leq 1$. Thus we have proved that $\sum_{m=1}^M \tilde{p}_m a_m^\circ \leq \alpha$. This means that \mathbf{a}° is feasible for the optimization problem (11). Combining (28) with (29), we have therefore proved that on the event \mathcal{E}_2 ,

$$\sum_{m=1}^M p_{P,m} \ell(\tilde{\pi}_m, \tilde{a}_m) \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}} + \left(1 + \frac{\sqrt{6}}{9} \right) \epsilon.$$

After combining with (26), we therefore have

$$\mathbb{E} \left[\text{leb}(\hat{C}_n(X_{n+1})) \right] \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}} + \left(1 + \frac{\sqrt{6}}{9} \right) \epsilon + \frac{\alpha}{n},$$

as long as $\epsilon \leq 1$. Furthermore, the assumption $\epsilon \leq 1$ ensures that $\log(4Mn/\alpha) \leq 2 \log n$, and so plugging in the definition of ϵ , we have proved that

$$\mathbb{E} \left[\text{leb}(\hat{C}_n(X_{n+1})) \right] \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\mathcal{X}_{1:M})}{\alpha}} + c' \sqrt{\frac{M \log n}{\alpha n}}$$

for a sufficiently large universal constant c' . If instead we have $\epsilon > 1$, then we have $\mathbb{E} \left[\text{leb}(\hat{C}_n(X_{n+1})) \right] \leq 1 \leq c'' \sqrt{\frac{M \log n}{\alpha n}}$ for a sufficiently large universal constant c'' . Taking $c = \max\{c', c''\}$, we have completed the proof of the theorem.

B.4. Proof of Corollary 1

To help with the proof, we begin by adapting our previous notation to the setting of a data-dependent partition. The partition $\hat{\mathcal{X}}_{1:M}^{\mathcal{R}}$ is a function of the first half of the data, i.e., data points $\{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor}$. Conditioning on these data points, we define

$$p_{P,m} = \mathbb{P}_{P_X} \left\{ X \in \hat{\mathcal{X}}_m^{\mathcal{R}} \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right\}$$

and

$$\pi_{P,m} = \mathbb{E}_{P_X} \left[\pi_P(X) \mid X \in \hat{\mathcal{X}}_m^{\mathcal{R}}; \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right],$$

where we should interpret this to mean that the partition $\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}$ is treated as fixed, and the probability and expectation are calculated with respect to an independent draw $X \sim P_X$. Similarly, we will write

$$\Delta(\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}) = \mathbb{E}_{P_X} \left[|\pi_P(X) - \pi_{P,m(X)}| \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right].$$

These quantities are now functions of the data points $\{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor}$.

Next we apply Theorem 3. Specifically, we will condition on the data points $\{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor}$ used to choose the partition so that the partition can be treated as fixed, and will apply Theorem 3 with $\lceil \frac{n}{2} \rceil \geq 2$ in place of n (i.e., we apply the theorem to data points $i = \lfloor \frac{n}{2} \rfloor + 1, \dots, n$ in place of $i = 1, \dots, n$). This proves that

$$\mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ \pi_P(X_{n+1}) \in \widehat{C}_n^{\mathcal{R}}(X_{n+1}) \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right\} \geq 1 - \alpha$$

and

$$\begin{aligned} \mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\text{leb}(\widehat{C}_n^{\mathcal{R}}(X_{n+1})) \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] \\ \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\Delta_P(\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}})}{\alpha}} + c\sqrt{\frac{M \log \lceil \frac{n}{2} \rceil}{\alpha \lceil \frac{n}{2} \rceil}}. \end{aligned}$$

Marginalizing over the data points $\{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor}$ used to choose the partition, we therefore have

$$\mathbb{P}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left\{ \pi_P(X_{n+1}) \in \widehat{C}_n^{\mathcal{R}}(X_{n+1}) \right\} \geq 1 - \alpha$$

and after applying Jensen's inequality,

$$\begin{aligned} \mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\text{leb}(\widehat{C}_n^{\mathcal{R}}(X_{n+1})) \right] \\ \leq L_\alpha(\Pi_P) + \sqrt{\frac{2\mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\Delta_P(\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}) \right]}{\alpha}} + c\sqrt{\frac{M \log \lceil \frac{n}{2} \rceil}{\alpha \lceil \frac{n}{2} \rceil}}. \end{aligned}$$

Note that these bounds hold for any distribution P . The first statement therefore immediately verifies that $\widehat{C}_n^{\mathcal{R}}$ satisfies the distribution-free coverage property (3). We now complete the proof of the bound on length. Write $\text{mid}_m = \frac{m(X) - \frac{1}{2}}{M}$, the midpoint of the range $[\frac{m-1}{M}, \frac{m}{M})$ of estimated probabilities that define the region $\widehat{\mathcal{X}}_m^{\mathcal{R}}$. We have

$$\begin{aligned} \Delta(\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}) &= \mathbb{E}_{P_X} \left[|\pi_P(X) - \pi_{P,m(X)}| \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] \\ &= \sum_{m=1}^M p_{P,m} \mathbb{E}_{P_X} \left[|\pi_P(X) - \pi_{P,m(X)}| \mid X \in \widehat{\mathcal{X}}_m^{\mathcal{R}}; \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] \end{aligned}$$

$$\begin{aligned} &\leq \sum_{m=1}^M p_{P,m} \left(|\pi_{P,m} - \text{mid}_m| + \right. \\ &\quad \left. \mathbb{E}_{P_X} \left[|\pi_P(X) - \text{mid}_m| \mid X \in \widehat{\mathcal{X}}_m^{\mathcal{R}}; \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] \right) \\ &\leq 2 \sum_{m=1}^M p_{P,m} \mathbb{E}_{P_X} \left[|\pi_P(X) - \text{mid}_m| \mid X \in \widehat{\mathcal{X}}_m^{\mathcal{R}}; \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] \\ &= 2 \mathbb{E}_{P_X} \left[|\pi_P(X) - \text{mid}_{m(X)}| \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right], \end{aligned}$$

where the next-to-last step holds since, for each m ,

$$\begin{aligned} |\pi_{P,m} - \text{mid}_m| &= \left| \mathbb{E}_{P_X} \left[\pi_P(X) \mid X \in \widehat{\mathcal{X}}_m^{\mathcal{R}}; \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] - \text{mid}_m \right| \\ &\leq \mathbb{E}_{P_X} \left[|\pi_P(X) - \text{mid}_m| \mid X \in \widehat{\mathcal{X}}_m^{\mathcal{R}}; \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right], \end{aligned}$$

by Jensen's inequality. Next, by definition of the partition $\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}$, we have

$$|\pi_P(X) - \text{mid}_{m(X)}| \leq \left| \pi_P(X) - \widehat{\pi}_{\lfloor \frac{n}{2} \rfloor}^{\mathcal{R}}(X) \right| + \frac{1}{2M}$$

almost surely. Therefore,

$$\begin{aligned} \mathbb{E} \left[\Delta(\widehat{\mathcal{X}}_{1:M}^{\mathcal{R}}) \right] &\leq \mathbb{E} \left[\frac{1}{M} + 2 \mathbb{E}_{P_X} \left[\left| \pi_P(X) - \widehat{\pi}_{\lfloor \frac{n}{2} \rfloor}^{\mathcal{R}}(X) \right| \mid \{(X_i, Y_i)\}_{i=1, \dots, \lfloor \frac{n}{2} \rfloor} \right] \right] \\ &= \frac{1}{M} + 2 \Delta_{\lfloor \frac{n}{2} \rfloor, P}(\mathcal{R}). \end{aligned}$$

Combining everything, we have

$$\begin{aligned} &\mathbb{E}_{(X_i, Y_i) \stackrel{\text{iid}}{\sim} P} \left[\text{leb}(\widehat{C}_n^{\mathcal{R}}(X_{n+1})) \right] \\ &\leq L_\alpha(\Pi_P) + \sqrt{\frac{2 \left(\frac{1}{M} + 2 \Delta_{\lfloor \frac{n}{2} \rfloor, P}(\mathcal{R}) \right)}{\alpha}} + c \sqrt{\frac{M \log \lceil \frac{n}{2} \rceil}{\alpha \lceil \frac{n}{2} \rceil}}. \end{aligned}$$

Plugging in our choice of M , we have proved the desired bound when the universal constant c' is chosen appropriately.

B.5. Proofs of supporting lemmas

Proof of Lemma 6. Without loss of generality we can take $t \leq \frac{1}{2}$. Furthermore, if $t' > \frac{1}{2}$ then $\ell(t', a + r\delta) \leq \ell(\frac{1}{2}, a + r\delta)$, so it suffices to consider only the case where $t' \leq \frac{1}{2}$ as well. Finally, if $t' < t$ then $\ell(t', a + r\delta) \leq \ell(t, a + r\delta) \leq \ell(t, a)$ since $\ell(t, a)$ is nondecreasing in $t \in [0, \frac{1}{2}]$ and nonincreasing in a , so the result is trivial in this case. Therefore, from this point on we only need to prove the result for $0 \leq t \leq t' \leq \frac{1}{2}$.

Now we split into cases.

- If $t = a = 0$, then $\ell(t, a) = \ell(0, 0) = 0$ and

$$\ell(t', a + r\delta) \leq \frac{t'}{2(a + r\delta)} \leq \frac{2\delta^2}{2 \cdot r\delta} = \frac{\delta}{r},$$

where the first step holds since $\ell(t, a) \leq \frac{t}{2a}$ for all t, a .

- If $a \geq \frac{1}{2}$, then

$$\begin{aligned} \ell(t', a + \delta r) - \ell(t, a) - \frac{\delta}{r} &= 2(1 - a - \delta r)t' - 2(1 - a)t - \frac{\delta}{r} \\ &= 2(1 - a)(t' - t) - 2\delta r t' - \frac{\delta}{r} \\ &\leq 2(1 - a)(t' - t) - \delta\sqrt{8t'} \quad \text{since } 2xy \leq x^2 + y^2 \text{ for all } x, y \\ &\leq t' - t - \delta\sqrt{8t'} \quad \text{since } a \geq \frac{1}{2} \text{ implies } 2(1 - a) \leq 1 \\ &\leq 0, \end{aligned}$$

where the last step holds by assumption on $|t' - t|$.

- If $t \leq a < \frac{1}{2}$, then $\ell(t, a) = \frac{t}{2a}$, and we also know that $\ell(t', a + \delta r) \leq \frac{t'}{2(a + \delta r)}$ as before. Therefore,

$$\begin{aligned} \ell(t', a + \delta r) - \ell(t, a) &\leq \frac{t'}{2(a + \delta r)} - \frac{t}{2a} \\ &\leq \frac{t + \delta\sqrt{8t} + 2\delta^2}{2(a + \delta r)} - \frac{t}{2a} \quad \text{by assumption on } |t' - t| \\ &\leq \frac{t + \delta r \cdot \frac{t}{a} + \frac{\delta}{r} \cdot 2a + 2\delta^2}{2(a + \delta r)} - \frac{t}{2a} \quad \text{since } 2xy \leq x^2 + y^2 \text{ for all } x, y \\ &= \frac{(a + \delta r) \cdot \left(\frac{t}{a} + \frac{2\delta}{r}\right)}{2(a + \delta r)} - \frac{t}{2a} = \frac{\delta}{r}. \end{aligned}$$

- If $a < t$ and $a + \delta r < t'$, then

$$\begin{aligned} \ell(t', a + \delta r) - \ell(t, a) - \frac{\delta}{r} &= \left(1 - \frac{a + \delta r}{2t'}\right) - \left(1 - \frac{a}{2t}\right) - \frac{\delta}{r} \\ &= \frac{a \cdot (t' - t)}{2tt'} - \frac{\delta r + 2t'\delta/r}{2t'} \\ &\leq \frac{t' - t}{2t'} - \frac{\delta r + 2t'\delta/r}{2t'} \quad \text{since } a < t \\ &\leq \frac{t' - t - \delta\sqrt{8t'}}{2t'} \quad \text{since } 2xy \leq x^2 + y^2 \text{ for all } x, y \\ &\leq 0, \end{aligned}$$

where the last step holds by assumption on $|t' - t|$.

- If $a < t$ and $a + \delta r \geq t'$, then $\ell(t, a) = 1 - \frac{a}{2t} \geq \frac{1}{2}$ and $\ell(t', a + \delta r) \leq \frac{t'}{2(a + \delta r)} \leq \frac{1}{2}$, and so

$$\ell(t', a + \delta r) - \ell(t, a) - \frac{\delta}{r} \leq 0.$$

We have now verified that the bound holds in all cases, which completes the proof. \square

Proof of Lemma 7. Fix any m . We will show that the bounds hold for this choice of m with probability at least $1 - \frac{\alpha}{Mn}$, and then the lemma follows by applying the union bound.

The multiplicative Chernoff bound (McDiarmid, 1998, Theorem 2.3(b,c)) states that, for any integer $N \geq 1$ and any $t \in [0, 1]$, for a random variable $B \sim \text{Binomial}(N, t)$, it holds for all $\delta > 0$ that

$$\mathbb{P} \left\{ B \geq Nt + \sqrt{2Nt \log(1/\delta)} \right\} \geq 1 - \delta \quad (30)$$

and

$$\mathbb{P} \left\{ B \leq Nt - \sqrt{2Nt \log(1/\delta)} \vee 3 \log(1/\delta) \right\} \geq 1 - \delta. \quad (31)$$

Now fix any m . We will prove that the statements (24) and (25) hold with probability at least $1 - \frac{\alpha}{Mn}$ for this m .

First, we have $n\hat{p}_m \sim \text{Binomial}(n, p_{P,m})$, and therefore, applying (30) with $N = n$, $t = p_{P,m}$, and $\delta = \frac{\alpha}{4Mn}$, with probability at least $1 - \frac{\alpha}{4Mn}$ it holds that

$$\hat{p}_m \geq p_{P,m} - \sqrt{p_{P,m} \cdot \frac{2 \log(4Mn/\alpha)}{n}}. \quad (32)$$

Furthermore, applying (31) with $N = n$, $t = p_{P,m}$, and $\delta = \frac{\alpha}{4Mn}$, with probability at least $1 - \frac{\alpha}{4Mn}$ it holds that

$$\hat{p}_m \leq p_{P,m} + \sqrt{p_{P,m} \cdot \frac{3 \log(4Mn/\alpha)}{n}} \vee \frac{3 \log(4Mn/\alpha)}{n}. \quad (33)$$

Now assume that (32) and (33) both hold. We will show that this implies (24) for this m . First, by (33), we have

$$\begin{aligned} \tilde{p}_m &= \hat{p}_m + \sqrt{\hat{p}_m \cdot \frac{3 \log(4Mn/\alpha)}{n}} + \frac{3 \log(4Mn/\alpha)}{n} \\ &\leq p_{P,m} + \sqrt{p_{P,m} \cdot \frac{18 \log(4Mn/\alpha)}{n}} + \frac{12 \log(4Mn/\alpha)}{n}. \end{aligned}$$

Next, by (32), we have

$$p_{P,m} \leq \hat{p}_m + \sqrt{\hat{p}_m \cdot \frac{2 \log(4Mn/\alpha)}{n}} + \frac{2 \log(4Mn/\alpha)}{n} \leq \tilde{p}_m - \frac{p_{P,m}}{n-1}.$$

where the last step holds since $n \geq 2$ and so $\frac{\log(4Mn/\alpha)}{n} \geq \frac{\log(4n)}{n} \geq \frac{1}{n-1} \geq \frac{p_{P,m}}{n-1}$. Thus we have

$$p_{P,m} \leq \tilde{p}_m \left(1 - \frac{1}{n}\right).$$

This verifies that (24) holds for this value of m .

Now we turn to (25). We will condition on \hat{p}_m . We split into cases:

- First suppose $\pi_{P,m}$ satisfies

$$\pi_{P,m} + \sqrt{\pi_{P,m} \cdot \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}} \vee \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m} \leq \frac{1}{2}.$$

Then $\hat{p}_m > 0$, and conditional on \hat{p}_m , we have

$$n\hat{p}_m\hat{\pi}_m \sim \text{Binomial}(n\hat{p}_m, \pi_{P,m}).$$

Therefore, applying (30) with $N = n\hat{p}_m$, $t = \pi_{P,m}$, and $\delta = \frac{\alpha}{4Mn}$, with probability at least $1 - \frac{\alpha}{4Mn}$ it holds that

$$\hat{\pi}_m \geq \pi_{P,m} - \sqrt{\pi_{P,m} \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}}. \quad (34)$$

And, applying (31) with $N = n\hat{p}_m$, $t = \pi_{P,m}$, and $\delta = \frac{\alpha}{4Mn}$, with probability at least $1 - \frac{\alpha}{4Mn}$ it holds that

$$\hat{\pi}_m \leq \pi_{P,m} + \sqrt{\pi_{P,m} \cdot \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}} \vee \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}. \quad (35)$$

Now assume that (32), (34), and (35) all hold. By definition of this case, the bound (35) immediately implies that $\hat{\pi}_m \leq \frac{1}{2}$ and therefore by definition, $\tilde{\pi}_m \leq \frac{1}{2}$ also. Applying (34),

$$\pi_{P,m} \leq \hat{\pi}_m + \sqrt{\hat{\pi}_m \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}} + \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m},$$

and $\pi_{P,m} \leq \frac{1}{2}$ by definition of this case, so

$$\tilde{\pi}_m = \min \left\{ \frac{1}{2}, \hat{\pi}_m + \sqrt{\hat{\pi}_m \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}} + \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m} \right\} \geq \pi_{P,m}.$$

By applying (35), we furthermore have

$$\tilde{\pi}_m \leq \hat{\pi}_m + \sqrt{\hat{\pi}_m \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}} + \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}$$

$$\leq \pi_{P,m} + \sqrt{\pi_{P,m} \cdot \frac{12 \log(4Mn/\alpha)}{n\hat{p}_m}} + \frac{8 \log(4Mn/\alpha)}{n\hat{p}_m}.$$

Next, if $\hat{p}_m \geq \frac{2}{3}p_{P,m}$ then this last bound yields

$$\tilde{\pi}_m \leq \pi_{P,m} + \sqrt{\pi_{P,m} \cdot \frac{18 \log(4Mn/\alpha)}{np_{P,m}}} + \frac{12 \log(4Mn/\alpha)}{np_{P,m}},$$

while if instead $\hat{p}_m < \frac{2}{3}p_{P,m}$ then by (32) we can see that $p_{P,m} \leq \frac{18 \log(4Mn/\alpha)}{n}$, and so

$$\tilde{\pi}_m \leq \frac{1}{2} \leq \frac{9 \log(4Mn/\alpha)}{np_{P,m}}.$$

Either way, then, we have verified that the statement (25) holds, as desired.

- Next, suppose $\pi_{P,m}$ satisfies

$$\pi_{P,m} \leq \frac{1}{2} < \pi_{P,m} + \sqrt{\pi_{P,m} \cdot \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}} \vee \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}.$$

If $\hat{p}_m > 0$, then conditional on \hat{p}_m , we have

$$n\hat{p}_m(1 - \hat{\pi}_m) \sim \text{Binomial}(n\hat{p}_m, 1 - \pi_{P,m}).$$

Therefore, applying (30) with $N = n\hat{p}_m$, $t = 1 - \pi_{P,m}$, and $\delta = \frac{\alpha}{4Mn}$, with probability at least $1 - \frac{\alpha}{4Mn}$ it holds that

$$(1 - \hat{\pi}_m) \geq (1 - \pi_{P,m}) - \sqrt{(1 - \pi_{P,m}) \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}}. \quad (36)$$

Furthermore the bound (34) holds with probability at least $1 - \frac{\alpha}{4Mn}$ as above. (If instead $\hat{p}_m = 0$, then the bounds (34) and (36) hold trivially.) Now assume that (32), (34), and (36) all hold. If $\hat{\pi}_m \leq \frac{1}{2}$, then by definition we have $\tilde{\pi}_m \leq \frac{1}{2}$ also. Furthermore, applying (34) proves that $\pi_{P,m} \leq \tilde{\pi}_m$ exactly as for the previous case. If instead $\hat{\pi}_m > \frac{1}{2}$, then (36) implies that

$$\frac{1}{2} \leq (1 - \pi_{P,m}) \leq (1 - \hat{\pi}_m) + \sqrt{(1 - \hat{\pi}_m) \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}} + \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m},$$

which therefore means that

$$\tilde{\pi}_m = \max \left\{ \frac{1}{2}, \hat{\pi}_m - \sqrt{(1 - \hat{\pi}_m) \cdot \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m}} - \frac{2 \log(4Mn/\alpha)}{n\hat{p}_m} \right\} = \frac{1}{2}.$$

We have therefore established that $\pi_{P,m} \leq \tilde{\pi}_m \leq \frac{1}{2}$ under either scenario. Next,

$$|\tilde{\pi}_m - \pi_{P,m}| \leq \frac{1}{2} - \pi_{P,m} < \sqrt{\pi_{P,m} \cdot \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}} \vee \frac{3 \log(4Mn/\alpha)}{n\hat{p}_m}$$

by definition of this case. As for the previous case, we can therefore verify (25) by considering the two possibilities $\hat{p}_m \geq \frac{2}{3}p_{P,m}$ and $\hat{p}_m < \frac{2}{3}p_{P,m}$, and applying (32).

- The case that $\pi_{P,m} > \frac{1}{2}$ is treated analogously.

We have now verified that, for each m , with probability at least $1 - \frac{\alpha}{Mn}$, the bounds (24) and (25) both hold, which completes the proof of the lemma. \square

Acknowledgements

R.F.B. was supported by the National Science Foundation via grant DMS-1654076, and by the Office of Naval Research via grant N00014-20-1-2337. The author thanks Emmanuel Candès, Haoyang Liu, Aaditya Ramdas, and Ryan Tibshirani for inspiration and many helpful discussions.

References

- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 2006. [MR2378491](#)
- Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Predictive inference with the jackknife+. *arXiv preprint arXiv:1905.02928*, 2019.
- Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. The limits of distribution-free conditional predictive inference. *Information and Inference: A Journal of the IMA*, iaaa017, <https://doi.org/10.1093/imaiai/iaaa017>, 2020. *arXiv preprint arXiv:1903.04684*.
- Adam D. Bull and Richard Nickl. Adaptive confidence sets in L^2 . *Probability Theory and Related Fields*, 156(3-4):889–919, 2013. [MR3078289](#)
- T. Tony Cai and Mark G. Low. Adaptive confidence balls. *The Annals of Statistics*, 34(1):202–228, 2006. [MR2275240](#)
- T. Tony Cai, Mark Low, and Zongming Ma. Adaptive confidence bands for nonparametric regression functions. *Journal of the American Statistical Association*, 109(507):1054–1070, 2014. [MR3265680](#)
- Alexandra Carpentier. Testing the regularity of a smooth signal. *Bernoulli*, 21(1):465–488, 2015. [MR3322327](#)
- Christopher Genovese and Larry Wasserman. Adaptive confidence bands. *The Annals of Statistics*, 36(2):875–905, 2008. [MR2396818](#)
- Evarist Giné and Richard Nickl. Confidence bands in density estimation. *The Annals of Statistics*, 38(2):1122–1170, 2010. [MR2604707](#)
- Evarist Giné and Richard Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*, volume 40. Cambridge University Press, 2016. [MR3588285](#)
- Chirag Gupta, Aleksandr Podkopaev, and Aaditya Ramdas. Distribution-free binary classification: prediction sets, confidence intervals and calibration. *arXiv preprint arXiv:2006.10564* (2020).

- László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer Science & Business Media, 2006. [MR1920390](#)
- Peter Hall and Joel Horowitz. A simple bootstrap method for constructing nonparametric confidence bands for functions. *The Annals of Statistics*, 41(4):1892–1921, 2013. [MR3127852](#)
- Marc Hoffmann and Richard Nickl. On adaptive inference and confidence bands. *The Annals of Statistics*, 39(5):2383–2409, 2011. [MR2906872](#)
- Heinrich Jiang. Non-asymptotic uniform rates of consistency for k-NN regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3999–4006, 2019.
- Jing Lei. Classification with confidence. *Biometrika*, 101(4):755–769, 2014. [MR3286915](#)
- Jing Lei and Larry Wasserman. Distribution-free prediction bands for nonparametric regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):71–96, 2014. [MR3153934](#)
- Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018. [MR3862342](#)
- Ker-Chau Li. Honest confidence regions for nonparametric regression. *The Annals of Statistics*, 17(3):1001–1008, 1989. [MR1015135](#)
- Mark G. Low. On nonparametric confidence intervals. *The Annals of Statistics*, 25(6):2547–2554, 1997. [MR1604412](#)
- Colin McDiarmid. Concentration. In *Probabilistic Methods for Algorithmic Discrete Mathematics*, pages 195–248. Springer, 1998. [MR1678578](#)
- Sahand N. Negahban, Pradeep Ravikumar, Martin J. Wainwright, and Bin Yu. A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers. *Statistical Science*, 27(4):538–557, 2012. [MR3025133](#)
- Ilya Nourtdinov, Volodya Vovk, Michael Vyugin, and Alex Gammerman. Pattern recognition and density estimation under the general iid assumption. In *International Conference on Computational Learning Theory*, pages 337–353. Springer, 2001. [MR2042045](#)
- Harris Papadopoulos. Inductive conformal prediction: Theory and application to neural networks. In *Tools in Artificial Intelligence*. InTech, 2008.
- Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *European Conference on Machine Learning*, pages 345–356. Springer, 2002. [MR2050303](#)
- Dominique Picard and Karine Tribouley. Adaptive confidence interval for pointwise curve estimation. *Annals of Statistics*, pages 298–335, 2000. [MR1762913](#)
- Mauricio Sadinle, Jing Lei, and Larry Wasserman. Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association*, 114(525):223–234, 2019. [MR3941250](#)
- Botond Szabó, Aad W. Van Der Vaart, and J.H. van Zanten. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *The Annals of Statistics*, 43(4):1391–1428, 2015. [MR3357861](#)

- Vladimir Vovk. Conditional validity of inductive conformal predictors. In *Asian Conference on Machine Learning*, pages 475–490, 2012. [MR3080332](#)
- Vladimir Vovk and Ivan Petej. Venn–Abers predictors. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pages 829–838, 2014.
- Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer Science & Business Media, 2005. [MR2161220](#)
- Grace Wahba. Bayesian “confidence intervals” for the cross-validated smoothing spline. *Journal of the Royal Statistical Society: Series B (Methodological)*, 45(1):133–150, 1983. [MR0701084](#)