Cross-Validation Conformal Risk Control

Kfir M. Cohen, Sangwoo Park, Osvaldo Simeone

Shlomo Shamai (Shitz)

King's Communications, Learning, and Information Processing (KCLIP) lab Viterbi Faculty of Electrical and Computing Engineering Centre for Intelligent Information Processing Systems (CIIPS)

Department of Engineering, King's College London

Technion-Israel Institute of Technology

Abstract—Conformal risk control (CRC) is a recently proposed technique that applies post-hoc to a conventional point predictor to provide calibration guarantees. Generalizing conformal prediction (CP), with CRC, calibration is ensured for a set predictor that is extracted from the point predictor to control a risk function such as the probability of miscoverage or the false negative rate. The original CRC requires the available data set to be split between training and validation data sets. This can be problematic when data availability is limited, resulting in inefficient set predictors. In this paper, a novel CRC method is introduced that is based on cross-validation, rather than on validation as the original CRC. The proposed cross-validation CRC (CV-CRC) extends a version of the jackknife-minmax from CP to CRC, allowing for the control of a broader range of risk functions. CV-CRC is proved to offer theoretical guarantees on the average risk of the set predictor. Furthermore, numerical experiments show that CV-CRC can reduce the average set size with respect to CRC when the available data are limited.

I. INTRODUCTION

A. Context and Motivation

One of the key requirements for the application of artificial intelligence (AI) tools to risk-sensitive fields such as healthcare and engineering is the capacity of AI algorithms to quantify their uncertainty [1], [2]. This requires guarantees on the adherence of the "error bars" produced by the AI model to the true predictive uncertainty. The predictive uncertainty encompasses both the epistemic uncertainty caused by limited availability of data and the aleatoric uncertainty inherent in the randomness of data generation [3]. Without making strong assumptions on the data generation mechanism it is generally impossible to provide strict uncertainty quantification

The work of Kfir M. Cohen, Sangwoo Park and Osvaldo Simeone has been supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement No. 725731. The work of Osvaldo Simeone has also been supported by an Open Fellowship of the EPSRC with reference EP/W024101/1, by the European Union's Horizon Europe Project CENTRIC under Grant 101096379, and by Project REASON, a UK Government funded project under the Future Open Networks Research Challenge (FONRC) sponsored by the Department of Science Innovation and Technology (DSIT). The work of Shlomo Shamai has been supported by the German Research Foundation (DFG) via the German-Israeli Project Cooperation (DIP), under Project SH 1937/1-1. The authors acknowledge King's Computational Research, Engineering and Technology Environment (CREATE). Retrieved January 18, 2024, from https://doi.org/10. 18742/rnvf-m076.

KMC, SP and OS conceived the project; KMC and SP developed the theory with the supervision and guidance of OS; KMC performed the simulation; KMC, SP and OS prepared the manuscript; and SS reviewed the text and contributed to the vision of the paper. All authors discussed the results and contributed to the final manuscript.



Illustration of (top) the existing validation-based conformal risk Fig. 1. control (VB-CRC) [7]; and (bottom) the proposed method cross-validationbased conformal risk control (CV-CRC), which aims at reducing the predictive sets sizes by reusing the available data \mathcal{D} more efficiently.

guarantees for any input, but assumption-free guarantees can be established on average over validation and test data [4]. Conformal prediction (CP) [5], [6], and its extension conformal risk control (CRC) [7], are widely established methodologies for the evaluation of predictors with provable uncertainty quantification properties.

To elaborate, assume access to a data set \mathcal{D} of N pairs of examples consisting of input x and output y. Based on the data set \mathcal{D} and on a class of point predictors, CP and CRC produce a set predictor $\Gamma(x|\mathcal{D})$ mapping a test input x into a subset of the output space. The size of the set predictor $\Gamma(x|\mathcal{D})$ provides a measure of the uncertainty of the predictor for input x [6]. On average over the data set \mathcal{D} and over a test input-output pair (x, y), we wish to guarantee the calibration condition

$$\mathbb{E}_{\boldsymbol{\mathcal{D}},\mathbf{x},\mathbf{y}\sim p_0(\mathcal{D},x,y)}\left[\ell\left(\mathbf{y},\Gamma(\mathbf{x}|\boldsymbol{\mathcal{D}})\right)\right] \leq \alpha, \tag{1}$$

where boldface fonts denote random quantities, $\ell(\cdot, \cdot)$ is a loss measure, and α a user-specified maximum average loss level. In (1), under the joint distribution $p_0(\mathcal{D}, x, y)$, the examples in the data set \mathcal{D} and the test pair (\mathbf{x}, \mathbf{y}) are assumed to be independent identically distributed (i.i.d.), or, more generally *exchangeable*.

CRC can satisfy the requirement (1) for any user-specified target average loss level α , as long as the loss function is bounded and it decreases as the predicted set grows. Examples of such loss functions are the 0-1 miscoverage probability

$$\ell(y,\Gamma) = \mathbb{1}(y \notin \Gamma), \tag{2}$$

which returns 1 if the true label y is not in the set Γ and 0 otherwise, and the *false negative rate*, which returns the fraction of true values of y that are not included in set Γ for multi-label problems [7].

The requirement (1) can be always satisfied for such monotonic loss functions by returning as set predictor Γ the entire set of possible values for the output variable y. However, a set predictor is useful only as long as it is of moderate average size. The motivation of this work is to construct a set predictor that meets (1), while producing small predictive sets even in the presence of a limited data set \mathcal{D} .

B. State of the Art

CP addresses the design of set predictors satisfying the calibration condition (1) in the special case of the miscoverage loss (2) [5], [6], [8]. There are several variants of CP, including validation-based CP (VB-CP), *cross-validation-based CP (CV-CP)* [9], and full CP [5]. While full CP is considered to be impractical, requiring many rounds of retraining, VB-CP splits the data set into training and validation data sets, and it operates over a single round of training. However, the need to devote a separate data set for validation can significantly reduce the quality of the trained model, resulting in predictive sets of large sizes when data are limited [9].

CV-CP reduces the computational complexity as compared to full CP, while reducing the predicted set size as compared to VB-CP. This is done by partitioning the available data set into multiple folds, each acting as a validation data set for the model trained based on leave-fold-out data. At the cost of increasing the complexity, requiring as many training rounds as the number of folds, CV-CP was shown to produce important savings in terms of prediction set sizes [10], [11], [12].

Other extensions of CP include CP-aware training strategies [13], [14], prediction under distributional shifts [15], improvements in the training algorithms [16], [17], novel calibration metrics [18], [19], applications to engineering problems [10], [20], and online versions [21], [22] with applications [23], [24].

CRC generalizes CP to address the calibration criterion (1) for a wider class of risks, with the only constraints that the risk function be bounded and monotonic in the predicted set size [7], [22], [25], [26]. The original CRC is validation-based, and hence it may be referred to as VB-CRC for consistency with the terminology applied above for CP. Accordingly, it

relies on a split of the data set into training and validation sets, resulting in inefficient predictive sets when data are limited.

C. Main Contributions

In this paper, we introduce a novel version of CRC based on cross-validation. The proposed CV-CRC method generalizes CV-CP, supporting arbitrary bounded and monotonic risk functions. As we will demonstrate, the design and analysis of CV-CRC are non-trivial extensions of CV-CP, requiring new definitions and proof techniques.

The rest of the paper is organized as follows. Sec. II provides the necessary background, while CV-CRC is presented in Sec. III. Numerical experiments are reported in Sec. IV, and Sec. V draws some conclusions. All proofs are deferred to the supplementary material.

II. BACKGROUND

Consider N + 1 data points

$$\underbrace{(\mathbf{x}[1], \mathbf{y}[1])}_{=\mathbf{z}[1]}, \underbrace{(\mathbf{x}[2], \mathbf{y}[2])}_{=\mathbf{z}[2]}, \dots, \underbrace{\mathbf{x}[N+1], \mathbf{y}[N+1])}_{=\mathbf{z}[N+1]} (3)$$

over the sample space $\mathcal{X} \times \mathcal{Y}$ that are drawn according to an *exchangeable* joint distribution $p_0(\mathcal{D}, x, y)$ over index $i = 1, \ldots, N$. The first N data points constitute the data set $\mathcal{D} = \{z[i] = (x[i], y[i])\}_{i=1}^N$, while the last data point z[N+1] is the test pair, which is also denoted as z = (x, y). We fix a loss function $\ell : \mathcal{Y} \times 2^{\mathcal{Y}} \to \mathbb{R}$, which, given any label $y \in \mathcal{Y}$ and a predictive set $\Gamma \subseteq \mathcal{Y}$, returns a loss bounded as

$$b \le \ell(y, \Gamma) \le B \tag{4}$$

for some constants $B < \infty$ and $b \in \{-\infty\} \cup \mathbb{R}$. We further require that the loss is monotonic in the predictive set Γ in the sense that the following implication holds

$$\Gamma_1 \subseteq \Gamma_2 \quad \Rightarrow \quad \ell(y, \Gamma_1) \ge \ell(y, \Gamma_2) \quad \text{ for each } y \in \mathcal{Y}.$$
 (5)

Note that the 0-1 miscoverage loss (2) assumed by CP satisfies (4) with b = 0 and B = 1, and it also satisfies the implication (5).

For a given data set D, VB-CRC uses a two-step procedure to satisfy the constraint (1) for some target average loss α in the interval

$$b \leq \alpha \leq B. \tag{6}$$

To start, as illustrated in the top panel of Fig. 1, the available data set \mathcal{D} is split into N^{tr} examples forming the *training set* \mathcal{D}^{tr} and $N^{\text{val}} = N - N^{\text{tr}}$ points forming the *validation set* \mathcal{D}^{val} with $\mathcal{D} = \mathcal{D}^{\text{tr}} \cup \mathcal{D}^{\text{val}}$. In the first step of VB-CRC, a model is trained based on the training set \mathcal{D}^{tr} using any arbitrary scheme. Then, in the second step, VB-CRC determines a threshold $\lambda \in \mathbb{R}$ by using the validation data set \mathcal{D}^{val} . As explained next, the threshold λ dictates which labels $y \in \mathcal{Y}$ are to be included in the prediction set $\Gamma_{\lambda}(x|\mathcal{D}^{\text{tr}})$ for any test input x as follows.

A nonconformity (NC) score $NC((x, y)|\mathcal{D}^{tr})$ is selected that evaluates the loss of the trained predictor on a pair (x, y). Examples of NC scores include the residual between the label and a trained predictor for regression problems and the log-loss for classification problems [6], [27], [28]. With the given NC score, the set prediction is obtained as

$$\Gamma_{\lambda}(x|\mathcal{D}^{\mathrm{tr}}) = \left\{ y' \in \mathcal{Y} \middle| \mathrm{NC}((x,y')|\mathcal{D}^{\mathrm{tr}}) \le \lambda \right\},$$
(7)

thus including all labels $y' \in \mathcal{Y}$ with NC score smaller or equal to the threshold λ . By design, the set (7) satisfies the nesting property

$$\lambda_1 < \lambda_2 \quad \Rightarrow \quad \Gamma_{\lambda_1}(x|\mathcal{D}^{\mathrm{tr}}) \subseteq \Gamma_{\lambda_2}(x|\mathcal{D}^{\mathrm{tr}})$$
(8)

for any input x and data sets \mathcal{D}^{tr} .

We define the *risk* as the population, or test, loss of the predicted set (7) as

$$R(\lambda|\mathcal{D}^{\mathrm{tr}}) = \mathrm{E}_{\mathbf{x},\mathbf{y}\sim p_0(x,y)} \Big[\ell \big(\mathbf{y}, \Gamma_\lambda(\mathbf{x}|\mathcal{D}^{\mathrm{tr}}) \big) \Big].$$
(9)

Given the validation data set $\mathcal{D}^{\text{val}} = \{(x^{\text{val}}[i], y^{\text{val}}[i])\}_{i=1}^{N^{\text{val}}}$, the risk (9) can be estimated as

$$\hat{R}^{\text{val}}(\lambda|\mathcal{D}^{\text{tr}},\mathcal{D}^{\text{val}}) = \frac{1}{N^{\text{val}}+1} \left(\sum_{i=1}^{N^{\text{val}}} \ell(y^{\text{val}}[i],\Gamma_{\lambda}(x^{\text{val}}[i]|\mathcal{D}^{\text{tr}})) + B\right)$$
(10)

which is a function of the threshold λ . This corresponds to a regularized, biased, empirical estimate of the risk (9) that effectively adds an (N+1)-th dummy validation example with maximal loss B.

VB-CRC chooses the lowest threshold λ such that the estimate (10) is no larger than the target average risk α as in

$$\lambda^{\mathrm{VB}}(\mathcal{D}^{\mathrm{val}}|\mathcal{D}^{\mathrm{tr}}) = \inf_{\lambda} \left\{ \lambda \Big| \hat{R}^{\mathrm{val}}(\lambda|\mathcal{D}^{\mathrm{tr}}, \mathcal{D}^{\mathrm{val}}) \le \alpha \right\}.$$
(11)

With this threshold choice, as proven in [7], the set predictor (7) obtained via VB-CRC, i.e.,

$$\Gamma^{\rm VB}(x|\mathcal{D}^{\rm tr},\mathcal{D}^{\rm val}) = \Gamma_{\lambda^{\rm VB}(\mathcal{D}^{\rm val}|\mathcal{D}^{\rm tr})}(x|\mathcal{D}^{\rm tr})$$
(12)

ensures the desired condition (1). More precisely, the condition (1) holds for any fixed training set \mathcal{D}^{tr} , i.e., we have the inequality

$$\mathbb{E}_{\boldsymbol{\mathcal{D}}^{\text{val}},\mathbf{x},\mathbf{y}\sim p_0(\mathcal{D}^{\text{val}},x,y)}\left[\ell\left(\mathbf{y},\Gamma^{\text{VB}}(\mathbf{x}|\mathcal{D}^{\text{tr}},\boldsymbol{\mathcal{D}}^{\text{val}})\right)\right] \leq \alpha.$$
(13)

Furthermore, in order for (13) to hold, VB-CRC only requires the validation data \mathcal{D}^{val} and test pair (x, y) to be exchangeable.

III. CROSS-VALIDATION CONFORMAL RISK CONTROL

While VB-CRC reviewed in the previous section guarantees the average risk condition (13), splitting the available data set into training and validation sets may potentially lead to inefficient set predictors, having large predictive sets on average. In this section, we introduce the proposed CV-CRC scheme that aims at improving the efficiency of VB-CRC [7] via crossvalidation [9], while still guaranteeing condition (1).

To start, as illustrated in the bottom panel of Fig. 1, the available data set $\mathcal{D} = \{z[i]\}_{i=1}^{N}$ is partitioned using a fixed mapping into K folds $\mathcal{D} = \{\mathcal{D}_k\}_{k=1}^{K}$ of N/K-samples each, which is assumed to be an integer. We will write each k-th fold as $\mathcal{D}_k = \{(x_k[1], y_k[1]), \dots, (x_k[N/K], y_k[N/K])\}$, and we will denote the mapping of the *i*-th data point z[i] to its

fold index as $k[i] : \{1, ..., N\} \rightarrow \{1, ..., K\}$. Like VB-CRC, CV-CRC operates in two steps.

In the first step, for any k-th fold, a model is trained using the leave-fold-out training set $\mathcal{D}_{-k} = \mathcal{D} \setminus \mathcal{D}_k$ of N - N/Ksamples. Accordingly, unlike VB-CRC, K training rounds are required for CV-CRC. In the second step, as we will detail, CV-CRC determines a threshold λ to determine which values of the output y to include in the predicted set.

Given a threshold λ , CV-CRC produces the predictive set

$$\Gamma_{\lambda}^{\text{CV}}(x|\mathcal{D}) = \left\{ y' \in \mathcal{Y} \middle| \min_{k \in \{1,\dots,K\}} \left\{ \text{NC}((x,y')|\mathcal{D}_{-k}) \right\} \le \lambda \right\},$$
(14)

which includes all labels $y' \in \mathcal{Y}$ with minimum, i.e., best case, NC score across the K folds, that is not larger than λ .

To determine the threshold λ , CV-CRC estimates the population risk (9) using cross-validation as

$$\hat{R}^{\text{CV}}(\lambda|\mathcal{D}) = \frac{1}{K+1} \left(\sum_{k=1}^{K} \sum_{j=1}^{N/K} \ell(y_k[j], \Gamma_\lambda(x_k[j]|\mathcal{D}_{-k})) + B \right).$$
(15)

The cross-validation-based estimate (15) can be interpreted as the conventional cross-validation loss evaluated on an augmented data set

$$\mathcal{D}^{\text{aug}} = \left\{ \underbrace{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K}_{=\mathcal{D}}, \mathcal{D}_{\text{dummy}} \right\}, \tag{16}$$

with the first K folds being the available data set $\mathcal{D} = \{\mathcal{D}_1, \ldots, \mathcal{D}_K\}$, and the additional (K + 1)-th fold containing N/K dummy points with the maximal loss of B. In a manner similar to VB-CRC, the addition to dummy data points acts as a regularizer for the estimate (15), which is required to provide performance guarantees.

Finally, CV-CRC selects the threshold λ by imposing that the cross-validation based estimate (15) of the loss is no larger than the target average loss value α as in

$$\lambda^{\rm CV}(\mathcal{D}) = \inf_{\lambda} \Big\{ \lambda \Big| \hat{R}^{\rm CV}(\lambda | \mathcal{D}) \le \alpha \Big\}.$$
(17)

CV-CRC reduces to the jackknife-minmax scheme in [9] when evaluated with the miscoverage loss (2) in the special case of K = N folds.

Theorem 1. Fix any bounded and monotonic loss function $\ell(\cdot, \cdot)$ satisfying conditions (4) and (5), and any NC score $NC((x, y)|\mathcal{D}^{tr})$ that is permutation-invariant with respect to the ordering of the examples in the training set \mathcal{D}^{tr} . For any number of folds satisfying $K \geq B/(\alpha - b) - 1$, the CV-CRC predictive set $\Gamma^{CV}_{\lambda^{CV}(\mathcal{D})}(x|\mathcal{D})$ with (14) and (17) guarantees the condition

$$\mathbb{E}_{\boldsymbol{\mathcal{D}},\mathbf{x},\mathbf{y}\sim p_{0}(\boldsymbol{\mathcal{D}},x,y)}\left[\ell\left(\mathbf{y},\Gamma^{CV}(\mathbf{x}|\boldsymbol{\mathcal{D}})\right)\right]\leq\alpha.$$
(18)

The theorem thus confirms that CV-CRC meets the desired condition (1). In this regard, we note that, as in (1), the average loss in (18) includes averaging over the entire data set D, unlike the condition (13) satisfied by VB-CRC. Furthermore, Theorem 1 requires the NC score to be permutation-invariant with respect to the data points in the training set, which is not



Fig. 2. Empirical risk of VB-CRC and CV-CRC for the vector regression problem.

the case for VB-CRC. Permutation-invariance is also needed for CV-CP [9], as well as for full CP [5]. In practice, a permutation-invariant NC score can be obtained by implementing permutation-invariant training schemes such as full gradient descent, in which the final trained model does not depend on the ordering of the training data points.

IV. EXAMPLES

In this section, we numerically validate the proposed CV-CRC using two synthetic examples. The first is a vector regression problem, whereas the second concerns the problem of temporal point process prediction [29], [30]. Our code is publicly available¹.

A. Vector Regression

Inspired by the example in [9], we first investigate a vector regression problem in which the output variable $y = [y_1, \ldots, y_m]^\top$ is *m*-dimensional. The joint distribution of data set \mathcal{D} and test pair (x, y) is obtained as

$$p_0(\mathcal{D}, x, y) = \int p_0(\phi) \left(\prod_{i=1}^{N+1} p_0(x[i]) p_0(y[i]|x[i], \phi) \right) \mathrm{d}\phi,$$
(19)

where (x[N+1] = x, y[N+1] = y) is the test example, and we have the Gaussian distributions

$$p_0(x) = \mathcal{N}(x|0, d^{-1}I_d),$$
 (20a)

$$p_0(y|x,\phi) = \mathcal{N}(y|\phi^\top \cdot x, \beta_0^{-1}I_m), \qquad (20b)$$

while $p_0(\phi)$ is a mixture of Gaussians with means determined by an i.i.d. Bernoulli vector **b** as

$$p_0(\phi) = \mathbf{E}_{\mathbf{b} \stackrel{\text{iid}}{\sim} \text{Bern}(0.5)} \left[\mathcal{N}(\phi | \mu_0 \mathbf{b}, \gamma_0^{-1} I_d) \right].$$
(21)

We set $\mu_0 = 10$, $\gamma_0 = 1$, $\beta_0 = 4$, $d = |\mathcal{X}| = 50$, and $m = |\mathcal{Y}| = 30$. Note that the distribution (19) is exchangeable.

¹https://github.com/kclip/cvcrc



Fig. 3. Empirical inefficiency of VB-CRC and CV-CRC for the vector regression problem.

Using maximum-likelihood learning, given a training data set \mathcal{D}^{tr} , we obtain the model parameter $\phi_{\mathcal{D}^{\text{tr}}}^{\text{ML}}$ used for the linear prediction model $\hat{y}(x|\mathcal{D}^{\text{tr}}) = (\phi_{\mathcal{D}}^{\text{ML}})^{\top}x$ as $\phi_{\mathcal{D}^{\text{tr}}}^{\text{ML}} = X_{\mathcal{D}^{\text{tr}}}^{\dagger}Y_{\mathcal{D}^{\text{tr}}}$, where $(\cdot)^{\dagger}$ denotes the pseudo-inverse, $(\cdot)^{\top}$ denotes transpose, and the input and label data matrices $X_{\mathcal{D}} \in \mathbb{R}^{N \times d}$ and $Y_{\mathcal{D}} \in \mathbb{R}^{N \times m}$ have input $(x^{\text{tr}}[i])^{\top}$ and label $(y^{\text{tr}}[i])^{\top}$ as their *i*th rows, respectively.

The NC score is set to the maximum prediction residual across the m dimensions of the output variable y as

$$\mathrm{NC}((x,y)|\mathcal{D}^{\mathrm{tr}}) = 2 \left\| y - \hat{y}(x|\mathcal{D}^{\mathrm{tr}}) \right\|_{\infty},$$
(22)

where the infinity norm $\|\cdot\|_{\infty}$ returns the largest magnitude of its input vector. This results in predictive sets (12) and (14) with (17) in the form of $\Gamma = \Gamma_1 \times \cdots \times \Gamma_m$, with \times being the Cartesian product and

$$\Gamma_{j}^{\text{VB}} = \left\{ y_{j} \mid |y_{j} - [\hat{y}(x|\mathcal{D}^{\text{tr}})]_{j}| \le \lambda^{\text{VB}}(\mathcal{D}^{\text{val}}|\mathcal{D}^{\text{tr}})/2 \right\}$$
(23)

with $[\cdot]_j$ standing for the *j*th element of its argument for VB-CRC, and

$$\Gamma_j^{\text{CV}} = \bigcup_{k=1}^K \left\{ y_j \left| \left| y_j - [\hat{y}(x|\mathcal{D}_{-k})]_j \right| \le \lambda^{\text{CV}}(\mathcal{D})/2 \right\}$$
(24)

for CV-CRC. The loss function used in the risk (1) is defined as

$$\ell(y,\Gamma) = \frac{1}{m} \sum_{j=1}^{m} \mathbb{1}\left(y_j \notin \Gamma_j\right),\tag{25}$$

which evaluates the fraction of entries of vector y that are not included in the predictive set. This loss satisfies condition (4) with b = 0 and B = 1. Note that CP is not applicable to this loss, since it is different from (2).

Lastly, we define the *inefficiency* as the size of the predictive set evaluated as the average over all dimensions of the predictive intervals across the *m* dimensions of the output *y*, i.e., $\operatorname{ineff}(\Gamma) = \frac{1}{m} \sum_{j=1}^{m} |\Gamma_j|$.

For target risk $\alpha = 0.1$, the empirical risk and empirical inefficiency of $N^{\text{te}} = 200$ test covariate-output pairs, averaged



Fig. 4. Temporal point process prediction: After observing the past d times t_1, \ldots, t_n , a point process set predictor outputs predictive intervals $\Gamma_j(x|\mathcal{D})$ for each of the next m points with $j = 1, \ldots, m$.



Fig. 5. Empirical risk (top) and inefficiency (bottom) of VB-CRC and *N*-CV-CRC for the temporal point process prediction problem.

over 50 independent simulations, are shown in Fig. 2 and Fig. 3. Fig. 2, validates the theoretical result that CRC schemes satisfy condition (1). However, from Fig. 3, VB-CRC is observed to have a larger inefficiency than CV-CRC, particularly in the small data set size regime. Thus, CV-CRC uses data more efficiently, with K = 20 folds striking a good balance between inefficiency and computational complexity in this regime.

B. Temporal Point Process Prediction

A temporal process consists of a sequence of events at random times t_1, t_2, \ldots with $t_1 < t_2 < \ldots$ As illustrated in Fig. 4, given the past d events' timings $x = \{t_1, \ldots, t_d\}$, the goal is to output intervals $\Gamma_j(x|\mathcal{D})$ for each of the following m events with $j = 1, \ldots, m$. The loss function is defined as in (25). Data and test sequences of timings are generated following a self-exciting Hawkes process [31] with intensity function

$$\lambda(t|\mathcal{H}_t) = \mu + \sum_{i:t_i < t} \left(\alpha_1 \beta_1 e^{-\beta_1(t-t_i)} + \alpha_2 \beta_2 e^{-\beta_2(t-t_i)} \right),$$

with $\mu = 0.2$, $\alpha_1 = \alpha_2 = 0.4$, $\beta_1 = 1$ and $\beta_2 = 20$ [29]. The predictor is a recurrent neural network that outputs a predictive density function $p(t_{i+1}|t_1, \ldots, t_i, \phi_{\mathcal{D}^{\mathrm{tr}}})$ with trained parameter $\phi_{\mathcal{D}^{\mathrm{tr}}}$ [29]. The median $\hat{t}_{i+1}(t_1, \ldots, t_i, \phi_{\mathcal{D}^{\mathrm{tr}}})$ of the predictive distribution is used as the point estimate for the (i + 1)-th event. For i > d, estimates $\{\hat{t}_j\}_{j=d+1}^{i-1}$ are used in lieu of the correct timings in the point prediction.

VB-CRC (12) produces intervals

$$\Gamma_j^{\text{VB}} = \left\{ y_j \left| \left| y_j - \hat{t}_{d+j}(\mathcal{D}^{\text{tr}}) \right| \le \gamma^j \lambda^{\text{VB}}(\mathcal{D}^{\text{val}}|\mathcal{D}^{\text{tr}})/2 \right\}, \quad (26)$$

where multiplication by the interval common ratio $\gamma = 1.2$ increases the interval sizes for later predictions, and for the CV-CRC (14), we have

$$\Gamma_j^{\text{CV}} = \bigcup_{k=1}^K \left\{ y_j \, \Big| \, |y_j - \hat{t}_{d+j}(\mathcal{D}_{-k})| \le \gamma^j \lambda^{\text{CV}}(\mathcal{D})/2 \right\}.$$
(27)

We set the length of the observed sequence as d = 60, and predict the next m = 6 events. We allow one event on average to lie outside the predicted intervals, i.e., $\alpha = 1/6$. We average over 200 independent simulations with $N^{\text{te}} = 1000$ test points in each run.

The top panel of Fig. 5 illustrate the test risk (25) as function of data set size N, validating that both scheme attain risks lower than the desired level α . The bottom panel of the figure shows that CV-CRC with K = N reduces the average size of the predicted intervals.

V. CONCLUSION

In this paper, we have introduced a novel conformal risk control (CRC) scheme based on cross-validation, generalizing cross-validation CP to losses beyond miscoverage. The proposed CV-CRC was shown to provably control the average risk, with experiments demonstrating it to be more efficient than VB-CRC when the available data for training and calibration are scarce. Further work may consider using the jackknife+ of [9] instead of the jackknife-minmax for more efficient predictive sets; and extending the scheme to meta-learning [32].

REFERENCES

- [1] D. Tran, J. Z. Liu, M. W. Dusenberry, D. Phan, M. Collier, J. Ren, K. Han, Z. Wang, Z. E. Mariet, H. Hu, N. Band, T. G. J. Rudner, Z. Nado, J. van Amersfoort, A. Kirsch, R. Jenatton, N. Thain, E. K. Buchanan, K. P. Murphy, D. Sculley, Y. Gal, Z. Ghahramani, J. Snoek, and B. Lakshminarayanan, "Plex: Towards Reliability using Pretrained Large Model Extensions," in *First Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward at ICML 2022*, 2022.
- [2] B. Rajendran, O. Simeone, and B. M. Al-Hashimi, "Towards efficient and trustworthy ai through hardware-algorithm-communication co-design," *arXiv preprint arXiv:2309.15942*, 2023.
- [3] O. Simeone, *Machine Learning for Engineers*. Cambridge University Press, 2022.
- [4] J. Lei and L. Wasserman, "Distribution-Free Prediction bands for nonparametric Regression," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 76, no. 1, pp. 71–96, 2014.
- [5] V. Vovk, A. Gammerman, and G. Shafer, Algorithmic Learning in a Random World. Springer, 2005, springer, New York.
- [6] A. N. Angelopoulos and S. Bates, "Conformal Prediction: A Gentle Introduction," *Foundations and Trends® in Machine Learning*, vol. 16, no. 4, pp. 494–591, 2023.
- [7] A. N. Angelopoulos, S. Bates, A. Fisch, L. Lei, and T. Schuster, "Conformal Risk Control," in *The Twelfth International Conference on Learning Representations*, 2024.
- [8] M. Fontana, G. Zeni, and S. Vantini, "Conformal Prediction: A Unified Review of Theory and New Challenges," *Bernoulli*, vol. 29, no. 1, pp. 1 – 23, 2023.
- [9] R. F. Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani, "Predictive Inference with the Jackknife+," *The Annals of Statistics*, vol. 49, no. 1, pp. 486–507, 2021.
- [10] K. M. Cohen, S. Park, O. Simeone, and S. Shamai Shitz, "Calibrating AI Models for Wireless Communications via Conformal Prediction," *IEEE Transactions on Machine Learning in Communications and Networking*, vol. 1, pp. 296–312, 2023.
- [11] N. Deutschmann, M. Rigotti, and M. R. Martinez, "Adaptive Conformal Regression with Jackknife+ Rescaled Scores," arXiv preprint arXiv:2305.19901, 2023.
- [12] C. Gupta, A. K. Kuchibhotla, and A. Ramdas, "Nested Conformal Prediction and Quantile Out-Of-Bag Ensemble Methods," *Pattern Recognition*, vol. 127, p. 108496, 2022.
- [13] D. Stutz, K. D. Dvijotham, A. T. Cemgil, and A. Doucet, "Learning Optimal Conformal Classifiers," in *International Conference on Learning Representations*, 2021.
- [14] B.-S. Einbinder, Y. Romano, M. Sesia, and Y. Zhou, "Training Uncertainty-Aware Classifiers with Conformalized Deep Learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 22380– 22395, 2022.
- [15] R. J. Tibshirani, R. Foygel Barber, E. Candes, and A. Ramdas, "Conformal Prediction under Covariate Shift," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [16] Y. Yang and A. K. Kuchibhotla, "Finite-Sample Efficient Conformal Prediction," arXiv preprint arXiv:2104.13871, 2021.
- [17] A. Kumar, S. Sarawagi, and U. Jain, "Trainable Calibration Measures for Neural Networks from Kernel Mean Embeddings," in *International Conference on Machine Learning*. PMLR, 2018, pp. 2805–2814.

- [18] M. J. Holland, "Making Learning More Transparent using Conformalized Performance Prediction," arXiv preprint arXiv:2007.04486, 2020.
- [19] A. Perez-Lebel, M. Le Morvan, and G. Varoquaux, "Beyond Calibration: Estimating the Grouping Loss of Modern Neural Networks," in *ICLR* 2023–The Eleventh International Conference on Learning Representations, 2023.
- [20] S. Park and O. Simeone, "Quantum Conformal Prediction for Reliable Uncertainty Quantification in Quantum Machine Learning," *IEEE Transactions on Quantum Engineering*, no. 01, pp. 1–24, nov 2023.
- [21] I. Gibbs and E. Candes, "Adaptive Conformal Inference Under Distribution Shift," Advances in Neural Information Processing Systems, vol. 34, pp. 1660–1672, 2021.
- [22] S. Feldman, L. Ringel, S. Bates, and Y. Romano, "Achieving Risk Control in Online Learning Settings," *Transactions on Machine Learning Research*, 2023.
- [23] Y. Zhang, S. Park, and O. Simeone, "Bayesian Optimization with Formal Safety Guarantees via Online Conformal Prediction," *arXiv preprint* arXiv:2306.17815, 2023.
- [24] K. M. Cohen, S. Park, O. Simeone, P. Popovski, and S. Shamai, "Guaranteed dynamic scheduling of ultra-reliable low-latency traffic via conformal prediction," *IEEE Signal Processing Letters*, vol. 30, pp. 473– 477, 2023.
- [25] S. Bates, A. Angelopoulos, L. Lei, J. Malik, and M. Jordan, "Distribution-Free, Risk-Controlling Prediction Sets," *Journal of the ACM (JACM)*, vol. 68, no. 6, pp. 1–34, 2021.
- [26] A. N. Angelopoulos, S. Bates, E. J. Candès, M. I. Jordan, and L. Lei, "Learn Then Test: Calibrating Predictive Algorithms to Achieve Risk Control," arXiv preprint arXiv:2110.01052, 2021.
- [27] A. N. Angelopoulos, S. Bates, M. I. Jordan, and J. Malik, "Uncertainty Sets for Image Classifiers using Conformal Prediction," in 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021.
- [28] Y. Romano, E. Patterson, and E. Candes, "Conformalized Quantile Regression," Advances in Neural Information Processing Systems, vol. 32, 2019.
- [29] T. Omi, N. Ueda, and K. Aihara, "Fully Neural Network Based Model for General Temporal Point Processes," in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019, pp. 2122–2132.
- [30] M. Dubey, R. Palakkadavath, and P. Srijith, "Bayesian Neural Hawkes Process for Event Uncertainty Prediction," *International Journal of Data Science and Analytics*, pp. 1–15, 2023.
- [31] A. G. Hawkes, "Spectra of Some Self-Exciting and Mutually Exciting Point Processes," *Biometrika*, vol. 58, no. 1, pp. 83–90, 1971.
- [32] S. Park, K. M. Cohen, and O. Simeone, "Few-Shot Calibration of Set Predictors via Meta-Learned Cross-Validation-Based Conformal Prediction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, no. 01, pp. 1–13, oct 2023.
- [33] A. K. Kuchibhotla, "Exchangeability, Conformal Prediction, and Rank Tests," arXiv preprint arXiv:2005.06095, 2020.
- [34] A. Dean and J. Verducci, "Linear Transformations that Preserve Majorization, Schur Concavity, and Exchangeability," *Linear Algebra and Its Applications*, vol. 127, pp. 121–138, 1990.

Cross-Validation Conformal Risk Control: Supplementary Material

APPENDIX A

PROOF THAT VB-CRC ACHIEVES TARGET RISK

In this appendix, we prove condition (13) for VB-CRC. While this result was originally shown in [7], here we provide an equivelant proof that is more convenient to support the proof of Theorem 1 in Appendix C. We start by bounding the VB-CRC threshold (11) using the following steps

$$\lambda^{\mathrm{VB}}(\mathcal{D}^{\mathrm{val}}|\mathcal{D}^{\mathrm{tr}}) = \inf_{\lambda} \left\{ \lambda \left| \frac{1}{N^{\mathrm{val}}+1} \left(\sum_{i=1}^{N^{\mathrm{val}}} \ell(y^{\mathrm{val}}[i], \Gamma_{\lambda}(x^{\mathrm{val}}[i]|\mathcal{D}^{\mathrm{tr}})) + B \right) \leq \alpha \right\}$$

$$\geq \inf_{\lambda} \left\{ \lambda \left| \frac{1}{N^{\mathrm{val}}+1} \left(\sum_{i=1}^{N^{\mathrm{val}}} \ell(y^{\mathrm{val}}[i], \Gamma_{\lambda}(x^{\mathrm{val}}[i]|\mathcal{D}^{\mathrm{tr}})) + \ell(y, \Gamma_{\lambda}(x|\mathcal{D}^{\mathrm{tr}})) \right) \right\} \right\}$$

$$=: \lambda'(\mathcal{D}^{\mathrm{val}}, x, y|\mathcal{D}^{\mathrm{tr}}),$$

$$(28)$$

where the inequality in (28) follows from (4). The ground-truth risk averaged over test example (x, y) and validation set \mathcal{D}^{val} is upper bounded as

$$\mathbb{E}_{\boldsymbol{\mathcal{D}}^{\mathrm{val}},\mathbf{x},\mathbf{y}\sim p_{0}(\mathcal{D}^{\mathrm{val}},x,y)} \left[\ell(\mathbf{y},\Gamma_{\lambda^{\mathrm{VB}}(\boldsymbol{\mathcal{D}}^{\mathrm{val}}|\mathcal{D}^{\mathrm{tr}})}(\mathbf{x}|\mathcal{D}^{\mathrm{tr}})) \right] \leq \mathbb{E}_{\boldsymbol{\mathcal{D}}^{\mathrm{val}},\mathbf{x},\mathbf{y}\sim p_{0}(\mathcal{D}^{\mathrm{val}},x,y)} \left[\ell(\mathbf{y},\Gamma_{\lambda'(\boldsymbol{\mathcal{D}}^{\mathrm{val}},\mathbf{x},\mathbf{y}|\mathcal{D}^{\mathrm{tr}})}(\mathbf{x}|\mathcal{D}^{\mathrm{tr}})) \right]$$

$$\leq \alpha.$$
(29b)

where the first inequality (29a) follows the nesting property (8) given inequality (28). The second inequality (29a) is an application of the following lemma, whose proof is deferred to Appendix B.

Lemma 2. Let $\mathbf{v}_1, \ldots, \mathbf{v}_M$ be random variables with an exchangeable joint distribution such that the equation $\mathbb{P}(\frac{1}{M}\sum_{i=1}^{M} \mathbf{v}_i \leq \alpha) = 1$ holds. Then, we have the inequality $\mathbb{E}_{\mathbf{v}_{1:M} \sim p_0(v_{1:M})}[\mathbf{v}_m] \leq \alpha$ for all $m \in \{1, ..., M\}$.

To apply Lemma 2 in (29a), we define $M = N^{\text{val}} + 1$ variables by

$$\mathbf{v}_{i} = \begin{cases} \ell \left(\mathbf{y}^{\text{val}}[i], \Gamma_{\lambda'(\boldsymbol{\mathcal{D}}^{\text{val}}, \mathbf{x}, \mathbf{y} | \mathcal{D}^{\text{r}})}(\mathbf{x}^{\text{val}}[i] | \mathcal{D}^{\text{tr}}) \right) & i = 1, \dots, N^{\text{val}} \\ \ell \left(\mathbf{y}, \Gamma_{\lambda'(\boldsymbol{\mathcal{D}}^{\text{val}}, \mathbf{x}, \mathbf{y} | \mathcal{D}^{\text{r}})}(\mathbf{x} | \mathcal{D}^{\text{tr}}) & i = N^{\text{val}} + 1, \end{cases}$$
(30)

whose empirical average is, by (28), no greater than α . Furthermore, to comply with the technical conditions of Lemma 2, variables $\mathbf{v}_{1:M}$ need to be exchangeable. This is justified by the following lemma, which is a corollary of [33, Theorem 3] or [34, Theorem 4].

Lemma 3. Let $\mathbf{w}_1, \ldots, \mathbf{w}_M \in \mathcal{W}$ be a collection of exchangeable random vectors, $f : \mathcal{W} \to \mathbb{R}$ be a fixed mapping, and $g : \mathcal{W}^M \to \mathbb{R}$ be a fixed mapping that is permutation-invariant, i.e., oblivious to the ordering of its M input values. Then, the M random variables formed as $\mathbf{v}_1 = f(\mathbf{w}_1, g(\mathbf{w}_{1:M})), \ldots, \mathbf{v}_M = f(\mathbf{w}_M, g(\mathbf{w}_{1:M}))$ are exchangeable.

Lemma 3 implies the exchangeability of variables (30) by defining the $N^{\text{val}} + 1$ exchangeable vectors as

$$\mathbf{w}_{i} = \begin{cases} \mathbf{z}^{\text{val}}[i] & i = 1, \dots, N^{\text{val}} \\ (\mathbf{x}, \mathbf{y}) & i = N^{\text{val}} + 1; \end{cases}$$
(31)

the permutation invariant function is set as $g(\cdot) = \lambda'(\cdot | \mathcal{D}^{tr})$; the fixed mapping is

$$\mathbf{v}_{i} = f\left(\mathbf{w}_{i} = (\mathbf{x}_{i}, \mathbf{y}_{i}), g(\mathbf{w}_{1:M})\right) = \ell(\mathbf{y}_{i}, \Gamma_{g(\mathbf{w}_{1:M})}(\mathbf{x}_{i}|\mathcal{D}^{\mathrm{tr}}));$$
(32)

and we focus on the average risk of the last term, i.e., $m = M = N^{\text{val}} + 1$. This completes the proof of (13).

APPENDIX B

Proof of Lemma 2

In this appendix, we prove Lemma 2. To start, define a *bag* $u = (u_1, \ldots, u_M)$ of M elements $u_1 \ldots, u_M$ as a multiset, i.e., as an unordered list with allowed repetitions [5]. By definition, two bags u and v are equal if they contain the same elements, irrespective of the ordering of their identical items, which we write as $u \stackrel{\text{bag}}{=} v$. One can form a bag out of a random vector $\mathbf{v}_1, \ldots, \mathbf{v}_M \sim p_0(v_{1:M})$ by discarding the order of the items. Accordingly, the distribution of the bag \mathbf{u} is given by

$$p_0(u) = \mathbb{P}\Big(\langle \mathbf{v}_1, ..., \mathbf{v}_M \rangle \stackrel{\text{bag}}{=} \langle u_1, ..., u_M \rangle \Big) = \sum_{\pi \in \Pi_M} \mathbb{P}\big(\mathbf{v}_1 = u_{\pi(1)}, ..., \mathbf{v}_M = u_{\pi(M)}\big), \tag{33}$$

where the sum is over the set Π_M of all M! permutations. For example, three Bernoulli variables $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \underset{i.i.d.}{\sim} \text{Bern}(q)$ with parameter $q \in [0, 1]$ can constitute four different bags. In fact, bag $\mathbf{u} \stackrel{\text{bag}}{=} \langle \mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \rangle$ equals $\langle 0, 0, 0 \rangle$ with probability (w.p.) $(1-q)^3$, (0,0,1) w.p. $3(1-q)^2q$, (0,1,1) w.p. $3(1-q)q^2$, and (1,1,1) w.p. q^3 . With these definitions, we obtain the following chain of inequalities

$$\mathbf{E}_{\mathbf{v}_{1:M} \sim p_0(v_{1:M})}[\mathbf{v}_m] = \mathbf{E}_{\mathbf{u} \sim p_0(u)} \left[\mathbf{E}_{\mathbf{v}_{1:M} \sim p_0(v_{1:M}|u)} \left[\mathbf{v}_m \right] \left[\mathbf{v}_{1,...,\mathbf{v}_M} \right] \stackrel{\text{bag}}{=} \mathbf{u} \right]$$

$$(34a)$$

$$= \mathbf{E}_{\mathbf{u} \sim p_0(u)} \left[\frac{1}{M} \sum_{l=1}^{M} \mathbf{u}_l \right]$$
(34b)

$$= \mathbf{E}_{\mathbf{v}_{1:M} \sim p_0(v_{1:M})} \left[\mathbf{E}_{\mathbf{u} \sim p_0(u|v_{1:M})} \left[\frac{1}{M} \sum_{l=1}^{M} \mathbf{u}_l \middle| \mathbf{u} \stackrel{\text{bag}}{=} \langle \mathbf{v}_1, \dots, \mathbf{v}_M \rangle \right] \right]$$
(34c)

$$= \mathbf{E}_{\mathbf{v}_{1:M} \sim p_0(v_{1:M})} \left[\mathbf{E}_{\mathbf{u} \sim p_0(u|v_{1:M})} \left[\frac{1}{M} \sum_{l=1}^{M} \mathbf{v}_l \middle| \mathbf{u} \stackrel{\text{bag}}{=} \langle \mathbf{v}_1, \dots, \mathbf{v}_M \rangle \right] \right]$$
(34d)

$$= \operatorname{E}_{\mathbf{v}_{1:M} \sim p_0(v_{1:M})} \left[\frac{1}{M} \sum_{l=1}^{M} \mathbf{v}_l \right]$$
(34e)

$$\leq \alpha.$$
 (34f)

The inequalities of (34) are justified as follows: (34a) and (34c) stem from the law of iterated expectations over all possible bags of M items; (34b) arises from the fact that each item in the bag has an equal likelihood to be the realization of the m-th variable \mathbf{v}_m ; is again the law of iterated expectation with the reintroduction of the random vector; (34d) stems from the fact that if two bags have the same items, their sum is identical; (34e) leverages the fact that the bag given its random variables is a deterministically specified; and lastly, (34f) is by the assumption in Lemma 2. This concludes the proof of Lemma 2.

APPENDIX C

PROOF OF THEOREM 1

To prove Theorem 1, let us introduced an augmented data set, such that the last, (K+1)-th, fold $\mathcal{D}^{te} = \mathcal{D}_{K+1}$ is composed of N/K arbitrary test points

$$\tilde{\mathcal{D}} = \left\{ \underbrace{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K}_{=\mathcal{D}}, \underbrace{\mathcal{D}_{K+1}}_{=\mathcal{D}^{\text{te}}} \right\}$$
(35)

with the test point (x, y) included as the first point in the test set, i.e., $(x, y) = (x^{\text{te}}[1], y^{\text{te}}[1]) = (x_{K+1}[1], y_{K+1}[1])$. By construction, all N + N/K points in the augmented data set \tilde{D} are exchangeable and distributed according to joint distribution $p_0(\tilde{\mathcal{D}}) = p_0(\mathcal{D}, \mathcal{D}^{\text{te}})$. We denote the elements of the augmented set $\tilde{\mathcal{D}}$ in (35) as

$$(\tilde{x}_k[j], \tilde{y}_k[j]) = (x_k[j], y_k[j]) \text{ for } k \in \{1, \dots K\}$$
(36a)

$$(\tilde{x}_{K+1}[j], \tilde{y}_{K+1}[j]) = (x^{\text{te}}[j], y^{\text{te}}[j]).$$
(36b)

Note that the augmented set $\tilde{\mathcal{D}}$ in (35) is different than the augmented set using dummy points \mathcal{D}^{aug} (16). For a pair of folds indices $k', k \in \{1, \dots, K+1\}$ with $k \neq k'$, we also define the augmented leave-two-folds-out (L2O) set as the augmented set without the two indexed folds, i.e.,

$$\tilde{\mathcal{D}}_{-(k',k)} = \tilde{\mathcal{D}} \setminus \{\mathcal{D}_{k'}, \mathcal{D}_k\}.$$
(37)

As a special case, when one of the indices points to the (K + 1)-th fold, which is the test fold, the L2O reduces to the leave-one-out of the available data set $\tilde{D}_{-(K+1,k)} = D_{-k}$. For every fold within the augmented data set \tilde{D} (35), we evaluate the average L2O loss (37), minimized over the second fold index as

$$\hat{R}_{\text{L2O}}^{\text{CV}}(\lambda|\tilde{\mathcal{D}}) = \frac{1}{K+1} \sum_{k=1}^{K+1} \frac{K}{N} \sum_{j=1}^{N/K} \min_{k' \in \{1,\dots,K+1\} \setminus \{k\}} \left\{ \ell \big(\tilde{y}_k[j], \Gamma_\lambda(\tilde{x}_k[j]|\tilde{\mathcal{D}}_{-(k,k')}) \big) \right\}.$$
(38)

Finally, we define the L2O threshold as the minimal threshold value for which the estimated average L2O risk (38) is no larger than α , i.e.,

$$\lambda_{\rm L2O}^{\rm CV}(\tilde{\mathcal{D}}) = \inf_{\lambda} \Big\{ \lambda \Big| \hat{R}_{\rm L2O}^{\rm CV}(\lambda|\tilde{\mathcal{D}}) \le \alpha \Big\}.$$
(39)

Corollary 4. The L2O threshold $\lambda_{L2O}^{CV}(\tilde{D})$ in (39) is fold-permutation-invariant, i.e., for any of the (K+1)! possible fold-permutation mappings π , we have

$$\lambda_{L2O}^{CV}(\tilde{\mathcal{D}}) = \lambda_{L2O}^{CV}(\{\tilde{\mathcal{D}}_k\}_{k=1}^{K+1}) = \lambda_{L2O}^{CV}(\{\tilde{\mathcal{D}}_{\pi[k]}\}_{k=1}^{K+1}).$$
(40)

This is due to the commutative property of the outer fold-summation and of the inner, within-fold, summation in (38). **Lemma 5.** The L2O threshold $\lambda_{L2O}^{CV}(\tilde{D})$ in (39) lower bounds the K-CV-CRC threshold (17)

$$\lambda_{L20}^{CV}(\tilde{\mathcal{D}}) \le \lambda^{CV}(\mathcal{D}). \tag{41}$$

The proof of Lemma 5 is given in Appendix D.

We now define K + 1 random variables $\mathbf{v}_1, \ldots, \mathbf{v}_{K+1}$, whose randomness stems from their dependence on the augmented data set $\tilde{\boldsymbol{\mathcal{D}}}$. Each k-th random variable \mathbf{v}_k is the minimal leave-two-fold-out empirical risk averaged over the N/K examples in the validation fold $\boldsymbol{\mathcal{D}}_k$, i.e.,

$$\mathbf{v}_{k} = \frac{K}{N} \sum_{j=1}^{N/K} \min_{k' \in \{1,\dots,K+1\} \setminus \{k\}} \left\{ \ell \left(\tilde{\mathbf{y}}_{k}[j], \Gamma_{\lambda_{\text{L2O}}^{\text{CV}}(\boldsymbol{\tilde{\mathcal{D}}})}(\tilde{\mathbf{x}}_{k}[j] | \boldsymbol{\tilde{\mathcal{D}}}_{-(k',k)}) \right) \right\} \text{ for } k = 1,\dots,K+1.$$

$$(42)$$

The random variables $\{\mathbf{v}_1, \ldots, \mathbf{v}_{K+1}\} = \{v_1(\lambda, \tilde{\boldsymbol{\mathcal{D}}}), \ldots, v_{K+1}(\lambda, \tilde{\boldsymbol{\mathcal{D}}})\}$ are exchangeable for any fixed threshold due to the exchangeability of the folds in the augmented data set. Therefore, by Lemma 2, we have the inequality

$$\mathbb{E}_{\boldsymbol{\tilde{\mathcal{D}}} \sim p_0(\tilde{\mathcal{D}})} \left[v_{K+1}(\lambda_{L20}^{CV}(\boldsymbol{\tilde{\mathcal{D}}}), \boldsymbol{\tilde{\mathcal{D}}}) \right] \leq \alpha.$$
(43)

We are now ready to follow the steps

$$E_{\mathcal{D},\mathbf{x},\mathbf{y}\sim p_{0}(\mathcal{D},x,y)}\left[\ell\left(\mathbf{y},\Gamma^{CV}(\mathbf{x}|\mathcal{D})\right)\right]$$
$$= E_{\mathcal{D},\mathbf{x},\mathbf{y}\sim p_{0}(\mathcal{D},x,y)}\left[\ell\left(\mathbf{y},\bigcup_{k'=1}^{K}\Gamma_{\lambda^{CV}(\mathcal{D})}(\mathbf{x}|\mathcal{D}_{-k'})\right)\right]$$
(44a)

$$\leq \mathbb{E}_{\boldsymbol{\mathcal{D}},\mathbf{x},\mathbf{y}\sim p_{0}(\boldsymbol{\mathcal{D}},x,y)}\left[\min_{k'\in\{1,\ldots,K\}}\left\{\ell\left(\mathbf{y},\Gamma_{\lambda^{\mathrm{CV}}(\boldsymbol{\mathcal{D}})}(\mathbf{x}|\boldsymbol{\mathcal{D}}_{-k'})\right)\right\}\right]$$
(44b)

$$= \mathbb{E}_{\boldsymbol{\mathcal{D}},\boldsymbol{\mathcal{D}}^{\mathsf{te}} \sim p_{0}(\mathcal{D},\mathcal{D}^{\mathsf{te}})} \left[\min_{k' \in \{1,\dots,K\}} \left\{ \ell \left(\mathbf{y}^{\mathsf{te}}[1], \Gamma_{\lambda^{\mathsf{CV}}(\boldsymbol{\mathcal{D}})}(\mathbf{x}^{\mathsf{te}}[1] | \boldsymbol{\mathcal{D}}_{-k'}) \right) \right\} \right]$$
(44c)

$$= \mathbb{E}_{\boldsymbol{\mathcal{D}},\boldsymbol{\mathcal{D}}^{\mathsf{te}} \sim p_{0}(\mathcal{D},\mathcal{D}^{\mathsf{te}})} \left[\frac{K}{N} \sum_{j=1}^{N/K} \min_{k' \in \{1,\dots,K\}} \left\{ \ell \left(\mathbf{y}^{\mathsf{te}}[j], \Gamma_{\lambda^{\mathsf{CV}}(\boldsymbol{\mathcal{D}})}(\mathbf{x}^{\mathsf{te}}[j] | \boldsymbol{\mathcal{D}}_{-k'}) \right) \right\}$$
(44d)

$$= \operatorname{E}_{\tilde{\boldsymbol{\mathcal{D}}} \sim p_{0}(\tilde{\mathcal{D}})} \left[\min_{k' \in \{1, \dots, K\}} \left\{ \frac{K}{N} \sum_{j=1}^{N/K} \ell(\tilde{\mathbf{y}}_{K+1}[j], \Gamma_{\lambda^{\operatorname{CV}}(\boldsymbol{\mathcal{D}})}(\tilde{\mathbf{x}}_{K+1}[j] | \tilde{\boldsymbol{\mathcal{D}}}_{-(k', K+1)})) \right\} \right]$$
(44e)

$$\leq \mathrm{E}_{\tilde{\boldsymbol{\mathcal{D}}} \sim p_{0}(\tilde{\mathcal{D}})} \Big[\min_{k' \in \{1, \dots, K\}} \Big\{ \frac{K}{N} \sum_{j=1}^{N/K} \ell \big(\tilde{\mathbf{y}}_{K+1}[j], \Gamma_{\lambda_{\mathrm{L20}}^{\mathrm{CV}}(\tilde{\boldsymbol{\mathcal{D}}})}(\tilde{\mathbf{x}}_{K+1}[j] | \tilde{\boldsymbol{\mathcal{D}}}_{-(k', K+1)}) \big) \Big\} \Big]$$
(44f)

$$\leq \alpha,$$
 (44g)

where (44a) is a consequence of (14), which is equivalent to $\Gamma_{\lambda}^{CV}(x|\mathcal{D}) = \bigcup_{k=1}^{K} \Gamma_{\lambda}(x|\mathcal{D}_{-k})$; inequality (44b) is due to the nesting property (5) applied on a particular left-fold-out k' which is a subset of the union of all left-fold-out sets; (44d) leverages exchangeability as all test points have the same expected loss; (44e) uses the augmented data set notations (36b); inequality (44f) is an outcome of the nesting properties (8) and (5) with inequality (41); in inequality (44g), we have used (43), alongside Corollary 4, stating that the L2O threshold is fold-invariant. This completes the proof of Theorem 1.

APPENDIX D PROOF OF LEMMA 5

The proof of Lemma 5 stated in Appendix C follows the steps

$$\lambda_{\text{L2O}}^{\text{CV}}(\tilde{\mathcal{D}}) = \inf_{\lambda} \left\{ \lambda \left| \frac{1}{K+1} \sum_{k=1}^{K+1} \frac{K}{N} \sum_{j=1}^{N/K} \min_{k' \in \{1,\dots,K+1\} \setminus \{k\}} \left\{ \ell \left(\tilde{y}_k[j], \Gamma_\lambda(\tilde{x}_k[j] | \tilde{\mathcal{D}}_{-(k,k')}) \right) \right\} \le \alpha \right\}$$
(45a)

$$= \inf_{\lambda} \left\{ \lambda \left| \frac{1}{K+1} \left(\sum_{k=1}^{K} \frac{K}{N} \sum_{j=1}^{N/K} \min_{\substack{k' \in \{1, \dots, K+1\} \setminus \{k\}}} \left\{ \ell \left(\tilde{y}_{k}[j], \Gamma_{\lambda}(\tilde{x}_{k}[j] | \tilde{\mathcal{D}}_{-(k,k')}) \right) \right\} \right.$$

$$(45b)$$

$$+\frac{K}{N}\sum_{j=1}^{K}\min_{k'\in\{1,\dots,K\}}\left\{\ell\left(\tilde{y}_{K+1}[j],\Gamma_{\lambda}(\tilde{x}_{K+1}[j]|\tilde{\mathcal{D}}_{-(K+1,k')})\right)\right\}\right) \leq \alpha \right\}$$

$$\leq \inf_{\lambda}\left\{\lambda\left|\frac{1}{K+1}\left(\sum_{k=1}^{K}\frac{K}{N}\sum_{j=1}^{N/K}\min_{k'\in\{1,\dots,K+1\}\setminus\{k\}}\left\{\ell\left(\tilde{y}_{k}[j],\Gamma_{\lambda}(\tilde{x}_{k}[j]|\tilde{\mathcal{D}}_{-(k,k')})\right)\right\}+\frac{K}{N}\sum_{j=1}^{N/K}B\right)\leq\alpha\right\}$$
(45c)

$$\leq \inf_{\lambda} \left\{ \lambda \left| \frac{1}{K+1} \left(\sum_{k=1}^{K} \frac{K}{N} \sum_{j=1}^{N/K} \ell \left(\tilde{y}_{k}[j], \Gamma_{\lambda}(\tilde{x}_{k}[j] | \tilde{\mathcal{D}}_{-(k,K+1)}) \right) + B \right) \leq \alpha \right\}$$
(45d)

$$= \inf_{\lambda} \left\{ \lambda \left| \frac{1}{K+1} \left(\sum_{k=1}^{K} \frac{K}{N} \sum_{j=1}^{N/K} \ell(y_k[j], \Gamma_\lambda(x_k[j]|\mathcal{D}_{-k})) + B \right) \le \alpha \right\}$$

$$= \lambda^{CV}(\mathcal{D}),$$
(45f)

$$(45f)$$

where (45a) stems from the definition in (39); (45b) is obtained by decomposing the first sum into its first K summation terms, and by listing the last term, the (K + 1)-th, on its own; and (45f) follows the definition in (17). This completes the proof of Lemma 5.