Notes on Algorithms and Regret Guarantees for Online Conformal Prediction

Dawei Xie
daweixie2029@u.northwestern.edu

Abstract

Conformal prediction is a distribution-free uncertainty quantification framework with finite-sample guarantees. Recent advances in conformal prediction have extended its standard framework to online learning settings. This report surveys the advances in online conformal prediction algorithms and highlights shared algorithmic and problem structures with online learning frameworks, with a particular focus on presenting two recent algorithms that build upon the Adaptive Conformal Inference algorithm and examining their coverage and regret guarantees.

1 Introduction

There are many tasks where machine learning models are expected to output the most likely response while also offering some form of actionable uncertainty quantification. However, prior uncertainty estimation methods typically require additional assumptions about data distributions, e.g., Bayesian neural networks need to "fit" a distribution over network weights Neal [2012], Gal and Ghahramani [2016]. Ensemble methods Lakshminarayanan et al. [2017], Jain et al. [2020], on the other hand, aim to characterize variability across multiple models, but rely critically on the included models and struggle to provide guarantees. In light of these challenges, conformal prediction has emerged as a rigorous, model-agnostic, distribution-free approach for constructing prediction sets/intervals with finite-sample coverage guarantees.

The only requirement for implementing conformal prediction is exchangeability for the joint dataset consisting of the calibration data (discussed in Section 3) and the new data, which is a weaker condition than i.i.d., yet is typically violated in online settings. To address this issue, a line of effort has been put in the development of conformal prediction frameworks with distribution shifts. By utilizing online learning approaches, recent advances in conformal prediction have extended its standard framework to online algorithms while still achieving promising guarantees.

The goal of this report is to survey and analyze coverage and regret guarantees for online conformal prediction algorithms, focusing on presenting the *Strongly Adaptive Online Conformal Prediction* (SAOCP) algorithm from Bhatnagar et al. [2023] and the *Dynamically-Tuned Adaptive Conformal Inference* (DtACI) algorithm from Gibbs and Candès [2024]. We aim to establish their regret improvement in a relatively unified manner, including coverage properties and dynamic regret. This report is structured as follows. In Section 2, we review related work and introduce online conformal prediction as a line of effort that incorporates online learning to tackle distribution shifts; In Section 3, we present the background and problem formulation of conformal prediction and highlight ACI's algorithmic developments; we present two meta-algorithms in Section 4; Section 5 concludes the report with discussion and summary.

2 Related Work

Online conformal prediction is motivated by the issue of distribution shift, where a key assumption of conformal prediction, *exchangeability*, is violated. To address this problem, there are roughly two distinct lines of effort. A line of work introduces weighting schemes by exploiting additional knowledge about the distribution shift Tibshirani et al. [2019], Podkopaev and Ramdas [2021], Barber et al. [2023], referred to as "weighted" conformal prediction in Prinster et al. [2024]. This direction of work also includes Xu and Xie [2021], Prinster et al. [2022], Fannjiang et al. [2022], Yang et al. [2024].

Another line of work aims to learn adaptive thresholds that are adjusted upon observing the miscoverage of recent instances Gibbs and Candès [2021], Feldman et al. [2023], Bhatnagar et al. [2023], Gibbs and Candès [2024], referred to as "adaptive" conformal prediction in Prinster et al. [2024]. We focus on the latter. In particular, many works out of this category incorporate online learning to design algorithms with dynamic regret guarantees. Some works also aim for algorithms with optimal regret in both stochastic and adversarial environments. This line of effort also includes Zaffran et al. [2022], Bastani et al. [2022], Angelopoulos et al. [2023a].

3 Problem Formulation

3.1 Conformal Prediction

We start by presenting a general form of conformal prediction. While a canonical classification (or regression) task is to find the true label (or the real-valued response) $Y_{\text{test}} \in \mathcal{Y}$ for a new instance $X_{\text{test}} \in \mathcal{X}$ with some model, conformal prediction Vovk et al. [2005], Angelopoulos et al. [2023b] returns a prediction set (or interval) $\mathcal{C}\left(X_{\text{test}}\right)$ such that on average

$$\mathbb{P}\left(Y_{\text{test}} \in \mathcal{C}\left(X_{\text{test}}\right)\right) \geq 1-\alpha \tag{1}$$

holds, where α is the pre-specified error rate (tolerance level). In particular, conformal prediction defines a score function $S : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$, assuming access to a base model $\hat{f}$. For a given dataset $\{(X_i, Y_i)\}_{i=1}^n$, such scores $S_i = S\left(X_i, Y_i\right)$ are sometimes referred to *nonconformity scores*, which empirically reflect how *unlikely* the model considers a given value to be the true response. An example score function could be $S(x, y) = |y - \hat{f}\left(x\right)|$.

In terms of the specific algorithmic procedure, there are two versions of conformal prediction, *split* and *full*. Split conformal prediction picks the $1-\alpha$ quantile $Q_{1-\alpha}$ of the nonconformity scores $\{S_i\}_{i=1}^n$ computed with the given dataset $\{(X_i, Y_i)\}_{i=1}^n$ (the *calibration* dataset), and constructs the prediction set/interval for X_{n+1} as

$$\mathcal{C}\left(X_{n+1}\right) = \left\{y : S\left(X_{n+1}, y\right) \leq Q_{1-\alpha}\right\} \tag{2}$$

Full conformal requires refitting the model for each possible $y \in \mathcal{Y}$. This creates for each y a corresponding dataset $\{(X_i, Y_i)\}_{i=1}^n \cup \{(X_{n+1}, y)\}$. As a result, each y also corresponds to a $1-\alpha$ quantile $Q_{1-\alpha}^y$ of the nonconformity scores $\{S_i^y\}_{i=1}^n$ computed with its corresponding dataset. The prediction set/interval is then constructed as

$$\mathcal{C}\left(X_{n+1}\right) = \left\{y : S^y\left(X_{n+1}, y\right) \leq Q_{1-\alpha}^y\right\} \tag{3}$$

Irrespective of the variant of conformal prediction, if data in $\{(X_i, Y_i)\}_{i=1}^n \cup \{(X_{n+1}, Y_{n+1})\}$ are exchangeable[1], forming prediction sets as in 2 or 3 ensures that 1 holds marginally.

There is a tension in conformal prediction between *coverage* and *efficiency*. For example, the desideratum for conformal classification is to have the prediction set as small as possible while still maintaining Eq 1. On the other hand, The $1-\alpha$ guarantee Eq 1 can be trivially satisfied by outputting the entire label space $1-\alpha$ of the time, and outputting an empty set α of the time.

Online conformal prediction instead studies the setting where the data $(X_1, Y_1), \ldots, (X_T, Y_T)$ are coming in a sequential fashion, which violates the exchangeability assumption. Online conformal prediction typically maintains a growing calibration dataset by including each new instance from each round.

[1]Random variables $\{Z_1, \ldots, Z_n\}$ are exchangeable if their joint distribution is unchanged under permutations:

$$\left(Z_1, \ldots, Z_n\right) \stackrel{d}{=} \left(Z_{\sigma(1)}, \ldots, Z_{\sigma(n)}\right),$$

for any permutations σ. Note that i.i.d. is an example (and a special case) of exchangeability.

3.2 Adaptive Conformal Inference (ACI)

Here we review one of the earliest works, the *Adaptive Conformal Inference* (ACI, Gibbs and Candès [2021]), which maintains an adaptive tolerance level α_t that gets updated with step size γ as

$$\alpha_{t+1} = \alpha_t + \gamma \left(\alpha - \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t}\right]\right) \tag{4}$$

where α is the pre-specified tolerance level, and $\mathcal{C}_t^{\alpha_t}$ is the prediction set for X_t using tolerance level α_t. At each time step t, ACI examines the coverage of the prediction set $\mathcal{C}_t^{\alpha_t}$ after the true response Y_t is revealed, increases or decreases the tolerance level accordingly, and then add (X_t, Y_t) to the dataset. An improvement can be made to this update to provide a more holistic evaluation of coverage while placing more weight to more recent errors:

$$\alpha_{t+1} = \alpha_t + \gamma \left(\alpha - \sum_{\tau=1}^{t} w_\tau \mathbb{1}\left[Y_\tau \notin \mathcal{C}_\tau^{\alpha_\tau}\right]\right)$$

where $\mathcal{C}_\tau^{\alpha_\tau}$ is the corresponding prediction set constructed with tolerance level α_τ at time τ, and $\{w_\tau\}_{1\leq\tau\leq t} \subseteq [0,1]$ is a sequence of increasing weights with $\sum_{\tau=1}^{t} w_\tau = 1$. However, Gibbs and Candès [2021] found that this alternative produced almost the same result as Eq 4.

It is also noted that the simple online update Eq 4 can be viewed as a gradient descent step with regard to the following pinball loss:

$$\mathcal{L}\left(\alpha_t^*, \alpha_t\right) = \left(\alpha - \mathbb{1}\left[\alpha_t^* < \alpha_t\right]\right)\left(\alpha_t^* - \alpha_t\right) \tag{5}$$

where α_t^* is the "true" (highest) tolerance level such that its corresponding prediction set is the smallest set that still covers the true value:

$$\alpha_t^* := \sup\left\{\alpha : Y_t \in \mathcal{C}_t^{\alpha}\right\}$$

We can then rewrite the update as

$$\begin{aligned}
\alpha_{t+1} &= \alpha_t + \gamma \left(\alpha - \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t}\right]\right) \\
&= \alpha_t + \gamma \left(\alpha - \mathbb{1}\left[\alpha_t^* < \alpha_t\right]\right) \\
&= \alpha_t - \gamma \nabla_{\alpha_t} \mathcal{L}\left(\alpha_t^*, \alpha_t\right)
\end{aligned}$$

Aside: Eq 5 shares a similar quantile structure as in the newsvendor problem. Suppose $S_t := S(X_t, Y_t)$ denotes the true nonconformity score of instance (X_t, Y_t) and that $\widehat{S}_t$ is the estimated score used as a threshold to pick labels. We assume $S_t \sim \mathcal{S}$ where $\mathcal{S}$ is some unknown nonconformity score distribution that depends on the unknown data distribution ($S : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$). This can be illustrated as

$$\begin{aligned}
C\left(\widehat{S}_t\right) &= h\left(\widehat{S}_t - S_t\right)^+ + b\left(S_t - \widehat{S}_t\right)^+ \\
&= \underbrace{\alpha \cdot \left(\widehat{S}_t - S_t\right)^+}_{\text{over-coverage}} + \underbrace{(1-\alpha) \cdot \left(S_t - \widehat{S}_t\right)^+}_{\text{under-coverage}}
\end{aligned}$$

which aligns with how $\widehat{S}_t$ is picked as $\frac{b}{b+h}$-quantile of the nonconformity scores. However, this formalization is insufficient to establish sub-interval validity guarantees and also deviates from the true objective of online conformal prediction, as discussed in the next subsection.

3.3 Analyzing Online Conformal Prediction Algorithms

The basic ACI algorithm has bounded coverage deviation:

$$\left|\frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t}\right] - \alpha\right| \leq \frac{\max\left\{\alpha_1, 1-\alpha_1\right\} + \gamma}{T\gamma} \tag{6}$$

The proof of this coverage bound can be found in AppendixA.1.

To evaluate online conformal prediction algorithms, it would be insufficient to just establish the coverage error. We would be interested in tacking the tension between coverage and efficiency across arbitrary time intervals. Specifically, while the regret over the entire horizon can be defined as

$$\text{Regret}(T) := \frac{1}{T}\sum_{t=1}^{T}\mathcal{L}\left(\alpha_t^*, \alpha_t\right) - \inf_{\alpha^\star}\frac{1}{T}\sum_{t=1}^{T}\mathcal{L}\left(\alpha_t^*, \alpha^\star\right)$$

we aim to establish a dynamic regret for any interval $I = [m, n] \subseteq [T]$:

$$\text{Regret}(I) := \frac{1}{|I|}\sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*, \alpha_t\right) - \frac{1}{|I|}\sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*, \alpha_t^\star\right) \tag{7}$$

We remark that $\mathcal{L}\left(\alpha_t^*, \alpha_t\right)$ in Eq 5 is not the true objective. Recall that α_t^* is the highest tolerance level that leads to the smallest prediction set covering the true Y_t. But our goal is to find α_t that's close to $\alpha_t^\star$ where $\mathbb{P}\left(Y_t \in \mathcal{C}_t^{\alpha_t^\star}\right) = 1 - \alpha$. We would not be expecting to have α_t correctly include X_t in each t (that'd be over-covering) – we aim for $1 - \alpha$ instead.

We will be presenting a comparative analysis of two modified online conformal algorithms (concretely, two meta-algorithms that manage multiple experts) with their regrets. We focus on the SAOCP algorithm that utilizes *strongly adaptive regret minimization* and measures the worst-case regret over all intervals of fixed length Bhatnagar et al. [2023], and the DtACI algorithm that supports tunable step size in gradient descent update Gibbs and Candès [2024].

4 Meta-Algorithms

4.1 Dynamically-Tuned Adaptive Conformal Inference (DtACI)

One of the major issues of ACI is that the choice of the step size critically affects its performance. To mitigate this, DtACI employs the idea of maintaining a candidate set of step sizes and learning to steer the algorithm towards picking the best step size. This meta-algorithm manages multiple experts to run multiple ACI algorithms in parallel. Each expert i corresponds to an ACI algorithm with the step size γ_i. Like the online learning algorithms we have seen in class (e.g., Hedge, EXP3), DtACI adopts an exponential re-weighting scheme to update the weights and inform the choice of the step size.

We outlined DtACI in Algorithm 1. Given the target error rate α, any pretrained model $\hat{f}$, and parameters σ and η, we assume K experts with a candidate set of step sizes $\{\gamma_i\}_{1\leq i\leq K}$ that we want to explore, initialized tolerance levels: $\left\{\alpha_1^i\right\}_{1\leq i\leq K}$ (which can be initialized as α), and initialized weights: $\left\{w_1^i\right\}_{1\leq i\leq K}$ (each initialized as 1). The algorithm also maintains a calibration dataset $\mathcal{D}_t$ that will contain past observations $\{(X_i, Y_i)\}_{i=1}^{t-1}$.

4.1.1 Coverage Error

To establish a bound for how the the coverage error deviates from the nominal error rate α, we assume instead that parameters σ and η are replaced with time-indexed σ_t and η_t. By defining $\text{err}_t = \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t}\right]$, we present the long-term coverage error deviation in the following theorem.

Theorem 1. *Suppose that on iteration t, the parameters η and σ are replaced by values η_t and σ_t. Let $\gamma_{\min} := \min \gamma_i$ and $\gamma_{\max} := \max \gamma_i$, then*

$$\left|\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left[\text{err}_t\right] - \alpha\right| \leq \frac{1+2\gamma_{\max}}{T\gamma_{\min}} + \frac{\left(1+2\gamma_{\max}\right)^2}{\gamma_{\min}}\frac{1}{T}\sum_{t=1}^{T}\eta_t e^{\eta_t(1+2\gamma_{\max})} + 2\frac{1+\gamma_{\max}}{\gamma_{\min}}\frac{1}{T}\sum_{t=1}^{T}\sigma_t,$$

where the expectation is over the randomness in Algorithm 1. In particular, if $\lim_{t\to\infty}\eta_t = \lim_{t\to\infty}\sigma_t = 0$, then $\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\text{err}_t \overset{a.s.}{=} \alpha$.

The proof can be found in AppendixA.2. We would be able to see that if σ_t and η_t decay to 0, the coverage rate will converge to the nominal error rate almost surely.

Algorithm 1 Dynamically-Tuned Adaptive Conformal Inference (DtACI) Gibbs and Candès [2024]

Require: Assuming K experts; target error rate α; parameters σ and η; pretrained model $\hat{f}$;

corresponding step sizes $\{\gamma_i\}_{1\leq i\leq K}$
corresponding initialized tolerance levels: $\left\{\alpha_1^i\right\}_{1\leq i\leq K}$
corresponding initialized weights: $\left\{w_1^i\right\}_{1\leq i\leq K}$
// $\mathcal{D}_t$ will contain past observations $\{(X_i, Y_i)\}_{i=1}^{t-1}$
// the true tolerance level α_t^* can be computed when Y_t is revealed.
for $t = 1, 2, \cdots, T$ **do**
  Update probabilities $p_t^i := w_t^i / \sum_{1\leq j\leq k} w_t^j, \forall 1 \leq i \leq k$.
  Pick tolerance level $\alpha_t = \alpha_t^i$ for this round with probability p_t^i.
  Observe X_t, output the corresponding prediction set $\mathcal{C}_t^{\alpha_t}$ based on α_t, $\hat{f}$ and $\mathcal{D}_t$.
  Construct prediction sets $\mathcal{C}_t^{\alpha_t^i}$ for each α_t^i based on $\mathcal{D}_t$ as well.
  Reveal the true label Y_t; $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(X_t, Y_t)\}$
  for $i = 1, 2, \cdots, K$ **do**
    $\bar{w}_t^i \leftarrow w_t^i \cdot \exp\left(-\eta \cdot \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right)$
    $w_{t+1}^i \leftarrow (1-\sigma)\bar{w}_t^i + \frac{\sigma}{K}\sum_{j=1}^K \bar{w}_t^j$
    $\alpha_{t+1}^i = \alpha_t^i + \gamma_i\left(\alpha - \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t^i}\right]\right)$
  end for
end for

4.1.2 Dynamic Regret

We now establish the dynamic regret in the following theorem.

Theorem 2. *Let $\gamma_{\max} := \max \gamma_i$ and assume that $\gamma_1 < \gamma_2 < \cdots < \gamma_K$ with $\gamma_{i+1}/\gamma_i \leq 2$ for all $1 < i \leq K$. Assume additionally that $\gamma_K \geq \sqrt{1 + 1/|I|}$ and $\sigma \leq 1/2$. Then, for any interval $I = [m, n] \subseteq [T]$ and any sequence $\alpha_m^\star, \ldots, \alpha_n^\star \in [0, 1]$,*

$$\frac{1}{|I|}\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*, \alpha_t\right)\right] - \frac{1}{|I|}\sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*, \alpha_t^\star\right) \leq \frac{\log(K/\sigma) + 2\sigma|I|}{\eta|I|} + \frac{\eta}{|I|}\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*, \alpha_t\right)^2\right] + 4\left(1+\gamma_{\max}\right)^2 \max\left\{\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^\star - \alpha_{t-1}^\star\right| + 1}{|I|}}, \gamma_1\right\}$$

where the expectation is over the randomness in Algorithm 1.

The proof can be found in AppendixA.3. It's worth noting that we can tune parameters to get a simpler bound with the assumption that

$$\gamma_1 \leq \sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^\star - \alpha_{t-1}^\star\right| + 1}{|I|}}$$

which can be viewed as a quantification of the distribution shift's "size". Then by setting $\sigma = 1/(2|I|)$ and $\eta = \sqrt{\frac{\log(2K|I|)+1}{\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}(\alpha_t^*, \alpha_t)^2\right]}}$, we can obtain

$$\begin{aligned}\frac{1}{|I|}\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*, \alpha_t\right)\right] - \frac{1}{|I|}\sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*, \alpha_t^\star\right) \leq& 2\sqrt{\frac{\log(2K|I|)+1}{|I|}}\sqrt{\frac{1}{|I|}\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*, \alpha_t\right)^2\right]} \\ &+ 4\left(1+\gamma_{\max}\right)^2\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^\star - \alpha_{t-1}^\star\right| + 1}{|I|}} \\ =& O\left(\sqrt{\frac{\log(|I|)}{|I|}}\right) + O\left(\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^\star - \alpha_{t-1}^\star\right|}{|I|}}\right)\end{aligned}$$

4.2 Strongly Adaptive Online Conformal Prediction

We first present an equivalent framing of conformal regression which aims to construct a prediction interval around the base prediction $\hat{f}(X_{n+1})$ with half width r:

$$\mathcal{C}^r(X_{n+1}) := \left[\hat{f}(X_{n+1}) - r, \hat{f}(X_{n+1}) + r\right]$$

Therefore, r would be the "radius" that we want to predict based on historical radii of the calibration data (which are exactly the nonconformity scores). This notion is different from conformalized quantile regression Romano et al. [2019].

Similarly, we can derive the pinball loss function for online conformal regression as

$$\mathcal{L}(r_t^*, \hat{r}_t) = (\alpha - \mathbb{1}[r_t^* > \hat{r}_t])(\hat{r}_t - r_t^*)$$

where r_t^* is the "true" (smallest) radius such that its corresponding prediction interval is the shortest interval that still covers the true value:

$$r_t^* := \inf\{r : Y_t \in \mathcal{C}_t^r\}$$

Adapted from Jun et al. [2017], SAOCP is also a meta-algorithm that manages a set of experts. Unlike DtACI which works with a fixed set of experts, SAOCP maintains an active set of experts. In each round t, the algorithm will initialize a new expert $\mathcal{I}_t$ and assign it a "lifetime" $L(t)$ defined by Eq 8, indicating the duration of expert $\mathcal{I}_t$ being active. it is seen that there will be at most $g\lfloor\log_2 t\rfloor$ active experts in any round t. This design enables implicitly placing more weight on recent instances and initializes the newly instantiated expert with the latest predicted radius.

We outlined SAOCP in Algorithm 2. We mark that each expert should be viewed as designed to be Scale-Free OGD Orabona and Pál [2018], which decays its learning rate based on cumulative gradient norms. In each round, SAOCP outputs the predicted radius by aggregating active experts' predicted radii. After observing the true response, the algorithm updates each active expert's weight and radius.

4.2.1 Coverage Error

To establish the coverage error, we would need to consider a randomized SAOCP. Instead of aggregating experts' predicted radii, we pick an expert's predicted radius $\widehat{r}_t^i$ to be the prediction with probability p_i, just like how we pick an expert's tolerance level in Algorithm1. Then by defining $\mathrm{err}_t = \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\widehat{r}_t}\right]$, SAOCP achieves

$$\left|\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[\mathrm{err}_t] - \alpha\right| \leq O\left(\inf_{\beta}\left(T^{1/2-\beta} + T^{\beta-1}S_\beta(T)\right)\right)$$

where $S_\beta(T)$ measures the expert weights' smoothness and each individual expert's cumulative gradient norms. The formal proof and the detailed explanation for $S_\beta(T)$ can be found in Appendix B.4 of Bhatnagar et al. [2023]. With some additional assumption that, if there exists some $\beta \in (1/2, 1)$ and $\gamma < 1 - \beta$ such that $S_\beta(T) \leq \widetilde{O}(T^\gamma)$, then we can obtain a simpler bound

$$\left|\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[\mathrm{err}_t] - \alpha\right| \leq \widetilde{O}\left(T^{-\min\{1/2-\beta, \beta-1+\gamma\}}\right) = o_T(1)$$

4.2.2 Dynamic Regret

We now establish the dynamic regret in the following theorem.

Theorem 3. *For any interval $I = [m, n] \subseteq [T]$, SAOCP in Algorithm2 achieves*

$$\sum_{t=m}^{n}\mathcal{L}(r_t^*, \widehat{r}_t) - \min_{r_{m:n}^\star}\sum_{t=m}^{n}\mathcal{L}(r_t^*, r_t^\star) \leq \widetilde{O}\left(B\left[\left(\sum_{t=m+1}^{n}\left|r_t^* - r_{t-1}^*\right|\right)^{1/3}|I|^{2/3} + \sqrt{|I|}\right]\right)$$

Algorithm 2 Strongly Adaptive Online Conformal Prediction (SAOCP) Bhatnagar et al. [2023]

Require: Target error rate α; radius bound B; learning rate $\eta \leftarrow B/\sqrt{3}$, pretrained model $\hat{f}$.

// $\mathcal{D}_t$ will contain past observations $\{(X_i, Y_i)\}_{i=1}^{t-1}$
// the true radius r_t^* can be computed when Y_t is revealed.
for $t = 1, 2, \cdots, T$ **do**
  Initialize a new expert $\mathcal{I}_t$ with α, η, and the predicted $\widehat{r}_{t-1}$ from last round (if $\widehat{r}_{t-1}$ exists); initialize its weight $w_t^t \leftarrow 0$.
  Active set of experts: $\mathcal{A}_t = \{\mathcal{I}_i : t - L(i) < i \le t\}$, where $L(i)$ is defined as

$$L(i) := g \cdot \max_{n \in \mathbb{Z}} \{2^n : i \equiv 0 \bmod 2^n\} \tag{8}$$

  // $g \in \mathbb{Z}_{\ge 1}$ is a multiplier; at most $g \lfloor \log_2 t \rfloor$ experts are active in each round t.
  Set prior probability $\pi_i \propto i^{-2} \left(1 + \lfloor \log_2 i \rfloor\right)^{-1} \mathbb{1}\left[\mathcal{I}_i \in \mathcal{A}_t\right]$
  Compute $\widehat{p}_i = \pi_i \left[w_t^i\right]_+$ for all $\mathcal{I}_i \in \mathcal{A}_t$ and normalize them as p_i.
  Aggregate experts' radii: $\widehat{r}_t = \sum_{i \in \mathcal{A}_t} p_i \widehat{r}_t^i$
  Observe X_t, output the corresponding prediction interval $\mathcal{C}_t^{\widehat{r}_t}$ based on $\widehat{r}_t$, $\hat{f}$ and $\mathcal{D}_t$.
  Reveal the true response Y_t; $\mathcal{D}_{t+1} = \mathcal{D}_t \cup \{(X_t, Y_t)\}$
  for $\mathcal{I}_i \in \mathcal{A}_t$ **do**
    $\widehat{r}_{t+1}^i = \widehat{r}_t^i - \eta \cdot \frac{\nabla \mathcal{L}\left(r_t^*, \widehat{r}_t^i\right)}{\sqrt{\sum_{j=1}^t \left\|\nabla \mathcal{L}\left(r_j^*, \widehat{r}_j^i\right)\right\|_2^2}}$
    $g_t^i = \begin{cases} \frac{1}{B}\left(\mathcal{L}\left(r_t^*, \widehat{r}_t\right) - \mathcal{L}\left(r_t^*, \widehat{r}_t^i\right)\right) & w_{i,t} > 0 \\ \frac{1}{B}\left[\mathcal{L}\left(r_t^*, \widehat{r}_t\right) - \mathcal{L}\left(r_t^*, \widehat{r}_t^i\right)\right]_+ & w_{i,t} \le 0 \end{cases}$
    $w_{t+1}^i = \frac{1}{t-i+1}\left(\sum_{j=i}^t g_j^i\right)\left(1 + \sum_{j=i}^t w_j^i g_j^i\right)$
  end for
end for

The proof can be found in AppendixA.4. It's worth noting that this is an application of the bound derived in Zhang et al. [2018]. Dividing it by $|I|$ gives

$$\frac{1}{|I|} \sum_{t=m}^{n} \mathcal{L}\left(r_t^*, \widehat{r}_t\right) - \min_{r_{m:n}^\star} \frac{1}{|I|} \sum_{t=m}^{n} \mathcal{L}\left(r_t^*, r_t^\star\right) \le \widetilde{O}\left(B\left[\left(\frac{\sum_{t=m+1}^n \left|r_t^* - r_{t-1}^*\right|}{|I|}\right)^{1/3} + 1/\sqrt{|I|}\right]\right)$$

5 Discussion and Summary

We note that both of the two presented meta-algorithms achieve valid long-term coverage. In terms of the dynamic regret that we are more interested in, we find that converting radius r back to tolerance level α to remove the bound factor B in SAOCP's dynamic regret yields $\widetilde{O}\left(\left[\sum_{t=m+1}^n \left|\alpha_t^* - \alpha_{t-1}^*\right| / |I|\right]^{1/3} + |I|^{-1/2}\right)$, while the dynamic regret of DtACI implies $\widetilde{O}\left(\left[\sum_{t=m+1}^n \left|\alpha_t^\star - \alpha_{t-1}^\star\right| / |I|\right]^{1/2}\right)$. This suggests that with mild additional assumptions, DtACI achieves better theoretical performance. It is also noted in Bhatnagar et al. [2023] that an earlier version of Gibbs and Candès [2024] had already established a better dependence on the average path length, but it adapted poorly when choosing the step size simultaneously for all intervals. It is seen that the current version of DtACI has alleviated the reliance on the choice of learning rate.

For a more fine-grained comparison, a crucial aspect of is to evaluate the algorithms in real-world applications. Gibbs and Candès [2024] also compared two other works built on ACI, Online Expert Aggregation on ACI (AgACI) Zaffran et al. [2022], and MultiValid Predictor (MVP) Bastani et al. [2022]. Online conformal prediction algorithms are typically evaluated in time series forecasting tasks. Two experiments adopted in Gibbs and Candès [2024] are stock market and Covid-19 case counts. However, we observe that the discrepancy in performance for different algorithms can be subtle in some tasks, and some algorithms (e.g., MVP) are even worse in performance compared to the naive baseline that uses a fixed tolerance level. While conformal prediction is generally expected to work without any distributional assumption, many online conformal algorithms are implicitly relying on assumptions that are not fully explored in the theoretical development. We expect the future line of work on online conformal prediction to be informed by advances in online learning, and such algorithms to be designed to exploit problem structures.

References

Anastasios Angelopoulos, Emmanuel Candes, and Ryan J Tibshirani. Conformal pid control for time series prediction. *Advances in neural information processing systems*, 36:23047–23074, 2023a.

Anastasios N Angelopoulos, Stephen Bates, et al. Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning*, 16(4):494–591, 2023b.

Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023.

Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Practical adversarial multivalid conformal prediction. *Advances in neural information processing systems*, 35:29362–29373, 2022.

Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning. In *International Conference on Machine Learning*, pages 2337–2363. PMLR, 2023.

Clara Fannjiang, Stephen Bates, Anastasios N Angelopoulos, Jennifer Listgarten, and Michael I Jordan. Conformal prediction under feedback covariate shift for biomolecular design. *Proceedings of the National Academy of Sciences*, 119(43):e2204569119, 2022.

Shai Feldman, Liran Ringel, Stephen Bates, and Yaniv Romano. Achieving risk control in online learning settings. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL `https://openreview.net/forum?id=5Y04GWvoJu`.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.

Isaac Gibbs and Emmanuel J. Candès. Adaptive conformal inference under distribution shift. In *Advances in Neural Information Processing Systems*, volume 34, pages 1660–1672, 2021.

Isaac Gibbs and Emmanuel J. Candès. Conformal inference for online prediction with arbitrary distribution shifts. *Journal of Machine Learning Research*, 25(162):1–36, 2024.

Paula Gradu, Elad Hazan, and Edgar Minasyan. Adaptive regret for control of time-varying dynamics. In *Learning for Dynamics and Control Conference*, pages 560–572. PMLR, 2023.

Elad Hazan. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.

Siddhartha Jain, Ge Liu, Jonas Mueller, and David Gifford. Maximizing overall diversity for improved uncertainty estimates in deep ensembles. In *Proceedings of the AAAI conference on artificial intelligence*, pages 4264–4271, 2020.

Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Improved strongly adaptive online learning using coin betting. In *Artificial Intelligence and Statistics*, pages 943–951. PMLR, 2017.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.

Francesco Orabona and Dávid Pál. Scale-free online learning. *Theoretical Computer Science*, 716: 50–69, 2018.

Aleksandr Podkopaev and Aaditya Ramdas. Distribution-free uncertainty quantification for classification under label shift. In *Uncertainty in Artificial Intelligence*, pages 844–853. PMLR, 2021.

Drew Prinster, Anqi Liu, and Suchi Saria. Jaws: Auditing predictive uncertainty under covariate shift. *Advances in Neural Information Processing Systems*, 35:35907–35920, 2022.

Drew Prinster, Samuel Don Stanton, Anqi Liu, and Suchi Saria. Conformal validity guarantees exist for any data distribution (and how to find them). In *International Conference on Machine Learning*, volume 235, pages 41086–41118. PMLR, 21–27 Jul 2024.

Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. *Advances in neural information processing systems*, 32, 2019.

Ryan J. Tibshirani, Rina Foygel Barber, Emmanuel J. Candès, and Aaditya Ramdas. Conformal prediction under covariate shift. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*, volume 29. Springer, 2005.

Chen Xu and Yao Xie. Conformal prediction interval for dynamic time-series. In *International Conference on Machine Learning*, pages 11559–11569. PMLR, 2021.

Yachong Yang, Arun Kumar Kuchibhotla, and Eric Tchetgen Tchetgen. Doubly robust calibration of prediction sets under covariate shift. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(4):943–965, 2024.

Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pages 25834–25866. PMLR, 2022.

Lijun Zhang, Tianbao Yang, Zhi-Hua Zhou, et al. Dynamic regret of strongly adaptive methods. In *International conference on machine learning*, pages 5882–5891. PMLR, 2018.

A Proofs

Detailed proofs can be found in Gibbs and Candès [2021, 2024], Bhatnagar et al. [2023].

A.1 Proof of ACI Coverage Bound

Proof We first show that $\forall t \in \mathbb{N}$,

$$\alpha_t \in [-\gamma, 1+\gamma] \tag{9}$$

We prove by contradiction that there exists some $\inf_t \alpha_t < -\gamma$ (It would be identical to do contradiction by assuming $\sup_t \alpha_t > 1+\gamma$). Note that $\sup_t |\alpha_{t+1} - \alpha_t| = \sup_t \gamma \, |\alpha - \mathbb{1}[Y_t \notin \mathcal{C}_t^{\alpha_t}]| < \gamma$. It would indicate that there is some $t \in \mathbb{N}$ such that $\alpha_t < 0$ and $\alpha_{t+1} < \alpha_t$. However,

$$\alpha_t < 0 \Longrightarrow Q_t(1-\alpha_t) = \infty \Longrightarrow \alpha_{t+1} = \alpha_t + \gamma \left(\alpha - \mathbb{1}[Y_t \notin \mathcal{C}_t^{\alpha_t}]\right) \geq \alpha_t$$

since the $\mathcal{C}_t^{\alpha_t}$ would now be the entire label space that trivially covers Y_t. A contradiction is thus reached. That is, for any $\alpha_t < 0$, it holds that $\alpha_{t+1} \geq \alpha_t$, and the difference between any α_t and α_{t+1} is less than γ, so there would not be any $\alpha_t < -\gamma$ (and similarly, no $\alpha_t > 1+\gamma$).

By expanding Eq 4, we can write

$$\alpha_{T+1} = \alpha_1 + \sum_{t=1}^{T} \gamma \left(\alpha - \mathbb{1}[Y_t \notin \mathcal{C}_t^{\alpha_t}]\right)$$

where $\alpha_{T+1} \in [-\gamma, 1+\gamma]$ as we have just proved. Rearranging the terms yields the result. □

A.2 Proof of DtACI's Coverage Error (Theorem1)

Proof Gibbs and Candès [2024](Appendix C.6) Let $\tilde{\alpha}_t := \sum_i p_t^i \alpha_t^i / \gamma_i$. We observe that

$$\begin{aligned}
\tilde{\alpha}_t &= \sum_i \frac{p_t^i \left(\alpha_{t+1}^i - \gamma_i \left(\alpha - \mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t^i}\right]\right)\right)}{\gamma_i} \\
&= \sum_i \frac{p_t^i \alpha_{t+1}^i}{\gamma_i} + \sum_i p_t^i \left(\mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t^i}\right] - \alpha\right) \\
&= \tilde{\alpha}_{t+1} + \sum_i \frac{\left(p_t^i - p_{t+1}^i\right) \alpha_{t+1}^i}{\gamma_i} + \sum_i p_t^i \left(\mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t^i}\right] - \alpha\right)
\end{aligned}$$

Thus,

$$\mathbb{E}\left[\mathbb{1}\left[Y_t \notin \mathcal{C}_t^{\alpha_t}\right]\right] - \alpha = \tilde{\alpha}_t - \tilde{\alpha}_{t+1} + \sum_i \frac{\left(p_{t+1}^i - p_t^i\right) \alpha_{t+1}^i}{\gamma_i} \tag{10}$$

For ease of notation, let $W_t := \sum_i w_t^i$ and $\tilde{p}_{t+1}^i := \frac{p_t^i \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right)}{\sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)}$. By definition,

$$p_{t+1}^i = \frac{w_{t+1}^i}{\sum_{i'} w_{t+1}^{i'}} = (1 - \sigma_t) \tilde{p}_{t+1}^i + \frac{\sigma_t}{k}$$

Then

$$\begin{aligned}
\tilde{p}_{t+1}^i - p_t^i &= \frac{p_t^i \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right)}{\sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)} - p_t^i \\
&= p_t^i \frac{\exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right) - \sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)}{\sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)} \\
&= p_t^i \frac{\sum_{i'} p_t^{i'} \left(\exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right) - \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)\right)}{\sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)} \\
&= p_t^i \frac{\sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right) \left(\exp\left(\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right) - \eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right) - 1\right)}{\sum_{i'} p_t^{i'} \exp\left(-\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right)\right)} \\
&= p_t^i \sum_{i'} \tilde{p}_{t+1}^{i'} \left(\exp\left(\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right) - \eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right) - 1\right).
\end{aligned}$$

By Eq 9 we know that $\alpha_t^i \in [-\gamma_i, 1 + \gamma_i]$ and thus that $\left|\mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right) - \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right| \leq \max\{\alpha, 1 - \alpha\} \left|\alpha_t^{i'} - \alpha_t^i\right| \leq 1 + 2\gamma_{\max}$. By the mean value theorem,

$$\left|\exp\left(\eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^{i'}\right) - \eta_t \mathcal{L}\left(\alpha_t^*, \alpha_t^i\right)\right) - 1\right| \leq \eta_t \left(1 + 2\gamma_{\max}\right) \exp\left(\eta_t \left(1 + 2\gamma_{\max}\right)\right)$$

and also,

$$\left|\tilde{p}_{t+1}^i - p_t^i\right| \leq p_t^i \eta_t \left(1 + 2\gamma_{\max}\right) \exp\left(\eta_t \left(1 + 2\gamma_{\max}\right)\right)$$

Applying Eq 9 again yields

$$\begin{aligned}
\left|\sum_i \frac{\left(p_{t+1}^i - p_t^i\right) \alpha_{t+1}^i}{\gamma_i}\right| &\leq (1 - \sigma_t) \sum_i \left|\frac{\left(\tilde{p}_{t+1}^i - p_t^i\right) \alpha_{t+1}^i}{\gamma_i}\right| + \sigma_t \sum_i \left|\frac{\left(1/k - p_t^i\right) \alpha_{t+1}^i}{\gamma_i}\right| \\
&\leq \frac{\eta_t \left(1 + 2\gamma_{\max}\right)^2}{\gamma_{\min}} \exp\left(\eta_t \left(1 + 2\gamma_{\max}\right)\right) + 2\sigma_t \frac{1 + \gamma_{\max}}{\gamma_{\min}}
\end{aligned}$$

Summing over t in Eq10 gives the following inequality,

$$\left|\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left[\mathrm{err}_t\right]-\alpha\right| \leq \frac{\left|\tilde{\alpha}_1-\tilde{\alpha}_{T+1}\right|}{T}+\frac{\left(1+2\gamma_{\max}\right)^2}{\gamma_{\min}}\frac{1}{T}\sum_{t=1}^{T}\eta_t e^{\eta_t(1+2\gamma_{\max})}+2\frac{1+\gamma_{\max}}{\gamma_{\min}}\frac{1}{T}\sum_{t=1}^{T}\sigma_t$$

Using Eq 9 one more time leads to $\gamma_{\min}\tilde{\alpha}_t \in [-\gamma_{\max}, 1+\ \gamma_{\max}]$ and thus $\left|\tilde{\alpha}_1-\tilde{\alpha}_{T+1}\right| \leq \left(1+2\gamma_{\max}\right)/\gamma_{\min}$. Plugging this back to the previous expression gives the bound.

□

A.3 Proof of DtACI's Dynamic Regret (Theorem2)

Proof We need two lemmas that are adapted from Gradu et al. [2023] and Hazan [2019] respectively.

Lemma 4. *(Adapted from Lemma A.2 in Gradu et al. [2023]) Assume that $\sigma \leq 1/2$. Then, for any interval $I=[m,n] \subseteq [T]$ and any $1 \leq i \leq K$*

$$\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*,\alpha_t\right)\right] \leq \sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*,\alpha_t^i\right)+\eta\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*,\alpha_t\right)^2\right]+\frac{1}{\eta}(\log(K/\sigma)+|I|2\sigma),$$

where the expectation is over the randomness in Algorithm1 and the data $\alpha_1^*,\ldots,\alpha_T^*$ can be viewed as fixed.

Lemma 5. *(Theorem 10.1 of Hazan [2019]) For any fixed interval $I=[m,n]$, sequence $\alpha_m^\star,\ldots,\alpha_n^\star$, and $1 \leq i \leq K$,*

$$\sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*,\alpha_t^i\right)-\sum_{t=m}^{n}\mathcal{L}\left(\alpha_t^*,\alpha_t^\star\right) \leq \frac{3}{2\gamma_i}\left(1+\gamma_i\right)^2\left(\sum_{t=m+1}^{n}\left|\alpha_t^\star-\alpha_{t-1}^\star\right|+1\right)+\frac{1}{2}\gamma_i|I|.$$

Now fix any $i \in [K]$, we can obtain

$$\begin{aligned}
&\sum_{t=r}^{s}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*,\alpha_t\right)\right]-\sum_{t=r}^{s}\mathcal{L}\left(\alpha_t^*,\alpha_t^\star\right)\\
&=\left(\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*,\alpha_t\right)\right]-\sum_{t=r}^{s}\mathcal{L}\left(\alpha_t^*,\alpha_t^i\right)\right)+\left(\sum_{t=r}^{s}\mathcal{L}\left(\alpha_t^*,\alpha_t^i\right)-\sum_{t=r}^{s}\mathcal{L}\left(\alpha_t^*,\alpha_t^\star\right)\right)\\
&\leq \underbrace{\eta\sum_{t=m}^{n}\mathbb{E}\left[\mathcal{L}\left(\alpha_t^*,\alpha_t\right)^2\right]+\frac{1}{\eta}(\log(K/\sigma)+|I|2\sigma)}_{\text{Applying Lemma 4}}\\
&+\underbrace{\frac{3}{2\gamma_i}\left(1+\gamma_i\right)^2\left(\sum_{t=m+1}^{n}\left|\alpha_t^\star-\alpha_{t-1}^\star\right|+1\right)+\frac{1}{2}\gamma_i|I|}_{\text{Applying Lemma 5}}
\end{aligned}$$

There are two cases. If

$$\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^*-\alpha_{t-1}^*\right|+1}{|I|}} \geq \gamma_1 \tag{11}$$

then since $\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^*-\alpha_{t-1}^*\right|+1}{|I|}} \leq \sqrt{1+1/|I|} \leq \gamma_k$, we can find γ_i such that

$$\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^\star-\alpha_{t-1}^\star\right|+1}{|I|}} \leq \gamma_i \leq 2\sqrt{\frac{\sum_{t=m+1}^{n}\left|\alpha_t^\star-\alpha_{t-1}^\star\right|+1}{|I|}}$$

Plugging this back to the previous expression gives the desired result. Otherwise if 11 does not hold, plug in γ_1 for γ_i. □

A.4 Proof of SAOCP's Dynamic Regret (Theorem3)

Proof This bound is obtained by applying the dynamic regret bound in Zhang et al. [2018] (Corollary 5). Note that Algorithm2 additionally implies iterates $\widehat{s}_t \in [-\eta, D+\eta] \subset [-D, 2D]$. We can obtain

$$\sum_{t=m}^{n} \mathcal{L}\left(r_t^*, \widehat{r}_t\right) - \min_{r_{m:n}^\star} \sum_{t=m}^{n} \mathcal{L}\left(r_t^*, r_t^\star\right) \leq \left(B\left[\widetilde{V}_{[m,n]}^{1/3} |I|^{2/3} + \sqrt{|I|}\right]\right)$$

where

$$\widetilde{V}_{[m,n]} = \sum_{t=m+1}^{n} \sup_{r' \in [0,B]} \left|\mathcal{L}\left(r_t^*, r'\right) - \mathcal{L}\left(r_{t-1}^*, r'\right)\right| \leq \sum_{t=m+1}^{n} \left|r_t^* - r_{t-1}^*\right| = V_{[m,n]}$$

The inequality is due to the fact that $\left|\mathcal{L}\left(r_t^*, r'\right) - \mathcal{L}\left(r_{t-1}^*, r'\right)\right| \leq \left|r_t^* - r_{t-1}^*\right|$ by the 1-Lipschitzness of the quantile loss function. $\square$