

Minimum-Length Conformal Prediction Sets for Ordinal Classification

Zijian Zhang^{*1}, Xinyu Chen^{*1}, Yuanjie Shi¹, Liyuan Lillian Ma², Zifan Xu², Yan Yan^{†1}

¹School of Electrical Engineering and Computer Science, Washington State University

²Independent researcher

{zijian.zhang2, xinyu.chen1, yuanjie.shi, yan.yan1}@wsu.edu, {liyuanma2015, zfan0780}@gmail.com

Abstract

Ordinal classification has been widely applied in many high-stakes applications, e.g., medical imaging and diagnosis, where reliable uncertainty quantification (UQ) is essential for decision making. Conformal prediction (CP) is a general UQ framework that provides statistically valid guarantees, which is especially useful in practice. However, prior ordinal CP methods mainly focus on heuristic algorithms or restrictively require the underlying model to predict a unimodal distribution over ordinal labels. Consequently, they provide limited insight into *coverage–efficiency trade-offs*, or a *model-agnostic and distribution-free* nature favored by CP methods. To this end, we fill this gap by propose an ordinal-CP method that is model-agnostic and provides *instance-level optimal* prediction intervals. Specifically, we formulate conformal ordinal classification as a minimum-length covering problem at the instance level. To solve this problem, we develop a sliding-window algorithm that is optimal on each calibration data, with only a linear time complexity in K , the # of label candidates. The local optimality per instance further also improves predictive efficiency in expectation. Moreover, we propose a length-regularized variant that shrinks prediction set size while preserving coverage. Experiments on four benchmark datasets from diverse domains are conducted to demonstrate the significantly improved predictive efficiency of the proposed methods over baselines (by 15%↓ on average over four datasets).

Code - <https://github.com/xrty/OCP>

1 Introduction

Ordinal classification plays a significant role in a wide range of real-world applications, e.g., medical diagnosis (Albuquerque, Cruz, and Cardoso 2021), credit risk assessment (Kwon, Han, and Lee 1997; Fernandez-Navarro et al. 2013; Hirk, Hornik, and Vana 2019), and age estimation (Niu et al. 2016). Unlike standard classification tasks that treat labels as unordered categories, ordinal classification explicitly models the inherent order among discrete labels. On one hand, this ordinal structure allows models to produce more interpretable predictions by reflecting the natural ranking in the

output (Gutiérrez et al. 2015). On the other hand, when the labels from the underlying distribution exhibit a clear ordinal structure, ordinal classification allows us to leverage the ordering information over classes, leading to better theoretical guarantees and empirical performance, e.g., Fisher consistency (Pedregosa, Bach, and Gramfort 2017).

While ordinal classification is highly valuable in real-world tasks, its practical effectiveness is often limited by its ability of uncertainty quantification (UQ) (Cresswell et al. 2024; Straitouri et al. 2023). In the absence of uncertainty-aware mechanisms, ordinal models typically produce a single point prediction without indicating statistically calibrated confidence. This can be problematic in safety-critical applications, especially when making a decision. For instance, in medical imaging, a model might assign a patient to a specific disease severity level based on subtle features (Zhang et al. 2024). However, if the model is poorly calibrated or highly uncertain due to limited data or domain shifts, such a point prediction may be misleading (Guo et al. 2017). Without a principled way to express uncertainty, practitioners may either over-rely on the model or disregard its output altogether, undermining its role in decision support. Therefore, integrating UQ into ordinal classification is essential to ensure reliable and interpretable deployment in high-stakes scenarios.

Conformal Prediction (CP) (Vovk, Gammerman, and Saunders 1999; Vovk, Gammerman, and Shafer 2005; Shafer and Vovk 2008; Angelopoulos and Bates 2021) is a post-hoc framework that provides distribution-free prediction sets with guaranteed marginal coverage (aka. statistically valid coverage), regardless of the underlying model’s confidence. Instead of committing to a single label as prediction, CP constructs a set of plausible labels for each instance, offering a principled way to quantify uncertainty. It is especially valuable in ordinal classification, where prediction confidence may vary significantly across classes and misclassifications may carry asymmetric consequences.

However, despite several efforts, extending standard CP to ordinal classification remains highly non-trivial. For example, (Lu, Angelopoulos, and Pomerantz 2022) adapts the seminal adaptive prediction sets (APS) (Romano, Sesia, and Candes 2020), which is originally designed for standard classification, and proposes ordinal APS. Their method differs from APS in the way to construct the prediction set –

^{*}These authors contributed equally.

[†]Corresponding author

given an input, it starts from the most probable label, and iteratively expands the prediction set by including the more confident neighboring label (either to the left or right), until the cumulative confidence exceeds a threshold calibrated on held-out calibration data. Nevertheless, this greedy search algorithm is inherently heuristic: it does not provide optimality guarantees and can yield larger-than-necessary intervals in some cases (i.e., worse predictive efficiency).

Another line of work, COPOC (Dey, Merugu, and Kaveri 2023), enforces unimodality on the model’s predicted distribution by learning an auxiliary transformation sub-module. While the ordinal structure encourages contiguous prediction sets, it still relies on a unimodality-enforcing module, which reduces the method’s model-agnostic nature.

To address the two aforementioned challenges for ordinal conformal prediction, we propose a novel algorithm, minimum-length conformal prediction sets (min-CPS). Analogous to prior methods, min-CPS aims to minimize the expected prediction interval subject to the marginal coverage constraint. However, instead of relying on heuristic or assuming model-specific conditions, min-CPS introduces a computationally efficient sliding-window algorithm that guarantees an exact solution to the minimum-length prediction set for each individual instance. In addition, we propose a length-regularized variant, i.e., min-RCPS, which explicitly incorporates the ordinal structure to further enhance efficiency, especially under uncertain intervals. This variant trades off between coverage tightness and structural alignment, leading to more semantically coherent prediction sets.

The main contributions of this paper are summarized as:

- We propose a novel algorithm, minimum-length conformal prediction sets (min-CPS), along with a length-regularized variant (min-RCPS) that leverages ordinal structure to improve predictive efficiency by balancing coverage tightness and semantic coherence.
- We provide theoretical guarantees for min-CPS, proving its instance-level optimality of the constructed intervals and establishing conditions that ensure monotonic empirical coverage. We empirically verify the monotonicity condition of the empirical coverage, under which a valid coverage is guaranteed.
- We conduct comprehensive experiments on four benchmark datasets, demonstrating that min-CPS consistently outperforms existing conformal prediction methods for ordinal classification in both predictive efficiency and interval coherence (min-CPS 14%↓ and min-RCPS 15%↓ reduction in the average length of prediction intervals).

2 Related Work

Ordinal Classification, also known as ordinal regression, addresses problems where the labels have a natural ordering but unknown spacing between categories (Wang et al. 2025; Gutiérrez et al. 2015). Unlike traditional multiclass classification, ordinal methods aim to respect this inherent order to improve predictive performance and interpretability. It has been widely used in real-world applications, e.g., facial age estimation (Shin, Lee, and Kim 2022; Wang et al. 2023),

aesthetics assessment (He et al. 2022; Kong et al. 2016), diabetic retinal grading (Yu et al. 2024; Cheng et al. 2023) and monocular depth estimation (Shao et al. 2023).

Conformal Prediction (CP) equips any base learner with *finite-sample, distribution-free* guarantees on predictive uncertainty, and has drawn increasing attention recently. Classical “full” CP computes leave-one-out non-conformity scores and inverts a series of hypothesis tests to ensure marginal coverage $1 - \alpha$ on unseen data (Vovk, Gamerman, and Shafer 2005; Shafer and Vovk 2008). Some studies focus on cross-validation methods (Vovk 2015) and jackknife+ approach (Barber et al. 2021), while its split (inductive) variant attains near-linear inference by calibrating a single quantile on a held-out set (Lei et al. 2018; Romano, Sesia, and Candes 2020; Oliveira et al. 2024), in which calibration is completed on a held-out dataset. Building on this foundation, researchers have devised score functions that tighten prediction sets without sacrificing validity. These examples include Adaptive Prediction Sets (APS) (Romano, Sesia, and Candes 2020), regularized APS (Angelopoulos et al. 2020), SAPS (Huang et al. 2024), Clustered CP (Ding et al. 2024), PoT-CP (Huang et al. 2025), RC3P (Shi et al. 2024), etc. For regression tasks, Conformalized Quantile Regression delivers valid confidence intervals (Romano, Patterson, and Candes 2019).

Recent work further refines CP for modern deep-learning by information-theoretic objectives (Correia et al. 2024), self-calibrating Venn–Abers schemes (van der Laan and Alaa 2024), probabilistic relaxations toward conditional validity (Plassier et al. 2024), robustness (Gendler et al. 2021; Ghosh et al. 2023; Liu et al. 2024), and generic conditional validity (Feldman, Bates, and Romano 2021; Gibbs, Cherian, and Candès 2023; Vovk 2012; Kiyani, Pappas, and Hassani 2024). Domain-specific extensions now support large-language-model inference with or without logit access (Cherian, Gibbs, and Candès 2024; Su et al. 2024), graph neural networks (Zhang et al. 2025), vision model (Chen et al. 2025), and enhance driver perception in adverse conditions through augmented-reality CP (Doula, Mühlhäuser, and Guinea 2024). A recent trend also arises to integrate CP into training (Stutz et al. 2022; Shi et al. 2025) to improve the model uncertainty. Together, these advances demonstrate CP’s versatility and its central role in uncertainty-aware machine learning.

Ordinal CP. Standard CP methods do not exploit the ordered structure of labels, often yielding non-contiguous prediction sets that are ill-suited to ordinal tasks. *Ordinal Adaptive Prediction Sets* (Ordinal-APS) introduce a score that accumulates *contiguous* soft-max probabilities, guaranteeing coverage while respecting the class order (Lu, Angelopoulos, and Pomerantz 2022). Building on this idea, (Dey, Merugu, and Kaveri 2023) cast the problem under a unimodality assumption and prove tight bounds on the minimal contiguous set size. Beyond coverage, (Xu, Guo, and Wei 2023) control more general risk measures (e.g., expected loss) in the ordinal setting, while (Chakraborty et al. 2024) leverage multiple-testing theory to construct both contiguous and non-contiguous sets with family-wise error control.

3 Proposed Min-CPS and Min-RCPS

In this section, we introduce our proposed algorithms, minimum-length conformal prediction sets (min-CPS) and its length-regularized variant minimum-length regularized conformal prediction sets (min-RCPS). We begin by the key notations and definitions used throughout the paper in §3.1. Next, in §3.2, we formulate the core building block, minimum-length covering formulation in the instance level, which serves as the foundation for both methods. We then show in §3.3 how this local formulation leads to global optimality of min-CPS in the population level, yielding provably efficient conformal prediction sets. Finally, in §3.4, we describe the regularized variant min-RCPS.

3.1 Notations and Problem Setup

Notations. Suppose $X \in \mathcal{X}$ is an input from the input space \mathcal{X} , and $Y \in \mathcal{Y} = \{1, 2, \dots, K\}$ is the ordinal ground-truth label, where K is the number of candidate classes. Assume that (X, Y) is a data sample drawn from an underlying distribution \mathcal{P} defined on the joint space $\mathcal{X} \times \mathcal{Y}$. Let $f(X) : \mathcal{X} \rightarrow \Delta_+^K$ denote an underlying pretrained model that predicts confidence for each class label, where Δ_+^K is the $(K-1)$ -dimensional probability simplex, $f(X)_y$ represents the predicted confidence score of class $y \in \mathcal{Y}$. Define $D_{\text{cal}} = \{(X_i, Y_i)\}_{i=1}^n$ as the calibration set of n examples, while $D_{\text{test}} = \{(X_i, Y_i)\}_{i=n+1}^{n+m}$ denotes the test set of m samples, all being exchangeable. Define $\mathbb{1}[\cdot]$ as an indicator function. \mathbb{N} represents the set of natural numbers.

Conformal Prediction for Ordinal Classification. Analogous to standard CP methods, ordinal CP algorithms aim to guarantee the following marginal coverage:

$$\mathbb{P}_{XY}\{Y \in \widehat{C}_\tau(X)\} \geq 1 - \alpha, \quad (1)$$

$$\text{where } \widehat{C}_\tau(X) = \{y \in \mathcal{Y} : l(X; \tau) \leq y \leq u(X; \tau)\},$$

where $l, u : \mathcal{X} \times \mathbb{R} \rightarrow \mathcal{Y}$ are some functions that map the input X and a calibration threshold τ to lower and upper bounds of the prediction interval, respectively.

In the context of ordinal classification, it is natural to impose the additional structural constraint that the prediction set forms a continuous interval over the label space. This reflects the semantic ordering of the labels and enhances interpretability, especially in high-stakes applications such as risk stratification or severity grading. If we relax the structure and allow $\widehat{C} : \mathcal{X} \rightarrow \{0, 1\}^{|\mathcal{Y}|}$ to include any arbitrary subset of \mathcal{Y} (not necessarily a continuous interval), then the framework reduces to a standard CP for multiclass classification, where the label order is ignored.

One notable exception arises when the model’s predicted probability distribution is unimodal over \mathcal{Y} . In this case, the standard CP methods naturally produce a contiguous prediction set on any X , since they sequentially include class labels in descending order of confidence, resulting in a single and connected interval around the mode (the most confident label). This property serves as the key motivation behind (Dey, Merugu, and Kaveri 2023), which explicitly enforces unimodality on predicted probability, so that standard CP methods can be directly applied.

Typically, the lower and upper functions l and u satisfy the *nestedness* property (Vovk, Gammerman, and Shafer 2005; Shafer and Vovk 2008) s.t. (with nonconformity scores):

$$\tau_1 \leq \tau_2 \Rightarrow \widehat{C}_{\tau_2}(X) \subseteq \widehat{C}_{\tau_1}(X),$$

which means that smaller thresholds lead to wider (more conservative) prediction intervals. This monotonicity property is a standard requirement in CP (Vovk, Gammerman, and Shafer 2005; Shafer and Vovk 2008), and is essential for calibrating the prediction set size to achieve the desired marginal coverage.

Accordingly, prior ordinal CP methods (Lu, Angelopoulos, and Pomerantz 2022; Dey, Merugu, and Kaveri 2023) share the same objective: to determine the smallest (least conservative) threshold τ on the held-out calibration set D_{cal} that guarantees the desired coverage, thus formulated as:

$$\begin{aligned} \min \quad & \tau \\ \text{s.t.} \quad & \sum_{i=1}^n \mathbb{1}[l(X_i; \tau) \leq Y_i \leq u(X_i; \tau)] \geq (1 - \alpha)(n + 1), \end{aligned} \quad (2)$$

which can be interpreted as *finding the least conservative coverage (by maximizing τ) that guarantee $1 - \alpha$ coverage*. Under the condition of data exchangeability, any τ satisfying the above constraint yields a conformal predictor with valid marginal coverage at level $1 - \alpha$ as in (1).

3.2 Instance-Level Minimum-Length Covering

Although solving (2) ensures the valid marginal coverage, existing studies have not provided insight into predictive efficiency (e.g., the expected size of the prediction intervals). For example, (Lu, Angelopoulos, and Pomerantz 2022) proposes a heuristic procedure for searching the lower and upper bounds of the prediction interval. However, the optimality gap of this approach remains under-explored, and it is unclear how close the resulting prediction sets are to the minimum possible length (best possible predictive efficiency). Hence, *it is still unclear whether existing ordinal CP methods reach the optimal predictive efficiency, and if not, what algorithms can be used to attain the optimum*.

Minimum-Length Covering on Each Instance. For any input $X \in \mathcal{X}$ and threshold $\tau \in (0, 1)$, we cast the search for the shortest interval as the following minimum-length covering problem:

$$\begin{aligned} \min_{(l, u) \in \mathcal{U}(X; \tau)} \quad & \underbrace{\ell(l, u) \triangleq u - l}_{\text{interval length}}, \quad \text{s.t. } \tau \in (0, 1), \text{ and} \quad (3) \\ \mathcal{U}(X; \tau) = \quad & \left\{ (l, u) : \underbrace{\sum_{k=l}^u f(X)_k}_{\text{(i) covering prob.}} \geq \tau, \underbrace{l \leq \hat{y}^*(X) \leq u}_{\text{(ii) including anchor}} \right\}, \end{aligned}$$

where $\hat{y}^*(X)$ is the mode of the predicted distribution (i.e., the most confident predicted label), and $\mathcal{U}(X; \tau)$ denotes a feasible set which requires the interval $[l, u]$ to (i) cover sufficient probability mass at least τ , and (ii) include the anchor label $\hat{y}^*(X)$.

other words, to achieve an ϵ -optimal threshold τ , the time complexity for min-CPS is $O(\log(1/\epsilon)nK)$, which is practically efficient (ref. running time comparison in Table 3).

Theorem 2. (*Coverage guarantee of min-CPS*) Under the same radial monotonicity assumption of Lemma 1, the calibrated threshold τ determined by Algorithm 2 yields the $(1 - \alpha)$ marginal coverage guarantee as in (1).

Comparison Problem (5) with Problem (2). As aforementioned, the standard formulation of ordinal CP in (2) does not necessarily take the predictive efficiency into account of the objective. Consequently, there is no guarantee of close to smallest intervals and it is unclear how to approach the optimal predictive efficiency. In contrast, by min-CPS, we solve an effective problem in (5), which indeed minimizes the interval length on each instance, and guarantees the superior predictive efficiency of min-CPS in practice.

3.4 Length-Regularization Variant: Min-RCPS

In addition to min-CPS, in this subsection, we further consider the semantic coherence of prediction intervals and propose a length-aware regularization into min-CPS to improve the predictive efficiency. Our motivation stems from the intuition that examples with larger minimum interval length $\ell(l^*(X; \tau), u^*(X; \tau))$ should be more uncertain than those with smaller minimum interval length, so we need to penalize the data with larger length via reducing their cumulative probability by a certainty quantity depending on their length.

Formally, we redefine the cumulative probability used in the feasible constraint of the instance-level minimum-length covering problem in (3) as follows:

$$\mathcal{U}_\lambda(X; \tau) \triangleq \left\{ (l, u) : \underbrace{\sum_{k=l}^u f(X)_k}_{\text{length reg.}} - \lambda \cdot \ell(l, u) \geq \tau, \right. \\ \left. l \leq \hat{y}^*(X) \leq u \right\}, \quad (6)$$

where we simply penalize the cumulative probability by a linear term of $\ell(l, u)$ scaled by a hyper-parameter λ . Beyond this redefinition of the feasible set $\mathcal{U}_\lambda(X; \tau)$, we keep everything else of min-CPS as it is, including solving the instance-level problem (3), and the binary search for τ as in Algorithm 2. We refer this algorithmic procedure as minimum-length regularized conformal prediction sets (min-RCPS). Analogous to min-CPS in (5), we also derive the objective of min-RCPS by using \mathcal{U}_λ of (6) as follows:

$$\min \tau \quad \text{s.t.} \quad (7) \\ \sum_{i=1}^n \mathbb{1}[l^*(X_i; \tau) \leq Y_i \leq u^*(X_i; \tau)] \geq \lceil (1 - \alpha)(n + 1) \rceil, \\ (l_\lambda^*(X_i; \tau), u_\lambda^*(X_i; \tau)) \in \arg \min_{(l, u) \in \mathcal{U}_\lambda(X_i; \tau)} \ell(l, u), \forall i.$$

Notably, when $\lambda = 0$, the regularization term vanishes and min-RCPS reduces min-CPS. We highlight that min-RCPS does not break the assumption of data exchangeability, so we have the following result:

Corollary 1. (*Marginal coverage of min-RCPS*) Min-RCPS preserves exchangeability and follows the same split-

conformal calibration rule as min-CPS. Therefore, the calibrated threshold $\hat{\tau}$ ensures the standard marginal coverage guarantee as in (1).

4 Experiments

4.1 Experimental Setup

Overview. We conduct extensive experiments to demonstrate the effectiveness of our proposed method on several real-world datasets. Following the literature of CP, we specifically evaluate two crucial metrics of ordinal CP methods: coverage and prediction set size. We compare our results against the state-of-the-art baseline Ordinal APS (Lu, Angelopoulos, and Pomerantz 2022), a naive method (referred to as Naive CDF, which is first used in (Lu, Angelopoulos, and Pomerantz 2022)) and a recent method WCRC (Xu, Guo, and Wei 2023) as a baseline.

Datasets. We evaluate our method on four publicly-available benchmark datasets for ordinal classification, which demonstrate both the modality diversity and the real-world relevance of ordinal classification:

- **UTKFace (Kaggle version)** (Kang 2019): A large-scale facial age dataset with over 20,000 images. Age is an inherently *ordered* variable, making it an ideal test-bed for ordinal classification methods.
- **Avocado Price** (Kiggins 2018): Historical weekly avocado prices discretised into ordinal price bands. This tabular, time-series dataset represents economic or financial forecasting scenarios where the outcome (price level) follows a natural ranking. It tests whether our approach generalises beyond vision data to structured numerical features subject to temporal trends and market volatility.
- **Electric Motor Temperature** (Kirchgeßner 2018): Multivariate sensor readings for monitoring motor temperature, discretised into ordinal risk/temperature ranges. Industrial condition-monitoring is a safety-critical domain: well-calibrated uncertainty is crucial to avoid both false alarms and missed overheating events. The sequential nature and class imbalance pose additional challenges for ordinal coverage guarantees.
- **IMDB** (Rothe, Timofte, and Gool 2015): A large-scale facial age dataset containing over 500,000 images collected from IMDB and Wikipedia. The dataset provides a challenging test-bed for ordinal classification and uncertainty quantification methods. Its scale, diverse real-world face conditions, and long-tailed age distribution make it a widely adopted benchmark in recent works (Keramati, Meng, and Evans 2024; Dong et al. 2025).

Together, these datasets span *image*, *tabular-economic*, and *time-series sensor* modalities, enabling us to assess the robustness and general applicability of our conformal ordinal predictors across heterogeneous real-world tasks.

Baselines. We consider three leading ordinal conformal prediction methods as baselines:

- **Ordinal APS** (Lu, Angelopoulos, and Pomerantz 2022): A variant of APS specifically tailored for ordinal classification, ensuring contiguous intervals.

Dataset	Method	Coverage	Prediction Set Size
Temperature	Naive CDF	0.9002 ± 0.0041	41.2510 ± 0.3770
	Ordinal APS	0.9015 ± 0.0040	5.3645 ± 0.0609
	WCRC	0.9025 ± 0.0030	5.3538 ± 5.0566
	min-CPS	0.9017 ± 0.0048	5.0902 ± 0.0581
	min-RCPS	0.9021 ± 0.0034	4.9913 ± 0.0428
UTKFace	Naive CDF	0.9025 ± 0.0048	73.6699 ± 0.9021
	Ordinal APS	0.9021 ± 0.0071	32.0085 ± 0.3504
	WCRC	0.9008 ± 0.0046	31.7998 ± 18.3118
	min-CPS	0.9026 ± 0.0055	29.7669 ± 0.2332
	min-RCPS	0.9025 ± 0.0055	29.7661 ± 0.2094
Avocado Price	Naive CDF	0.9012 ± 0.0046	23.2533 ± 0.0693
	Ordinal APS	0.9024 ± 0.0030	15.1795 ± 0.0596
	WCRC	0.9014 ± 0.0043	15.1587 ± 4.9565
	min-CPS	0.9034 ± 0.0057	9.1179 ± 0.0964
	min-RCPS	0.9033 ± 0.0054	8.9235 ± 0.0510
IMDB	Naive CDF	0.8984 ± 0.0062	37.3994 ± 0.6318
	Ordinal APS	0.9022 ± 0.0064	29.0753 ± 0.3479
	WCRC	0.9008 ± 0.0000	29.1937 ± 0.0122
	min-CPS	0.9022 ± 0.0055	28.1761 ± 0.3339
	min-RCPS	0.9019 ± 0.0072	28.2635 ± 0.3791

Table 1: Coverage and prediction set size ($\alpha = 0.1$). Bold numbers indicate the smallest prediction set size, while all methods guarantees at least $1 - \alpha$ coverage rate. On Temperature, UTKFace, Avocado Price and IMDB, comparison with the best baseline (Ordinal APS), min-CPS reduces the prediction set size by 5.11%↓, 7%↓, 39.93%↓, 3.49%↓, while min-RCPS reduces it by 6.96%↓, 7.01%↓, 41.21%↓, 3.16%↓. Average over all datasets, min-CPS and min-RCPS reduce the prediction set size by 14%↓ and 15%↓.

- **Naive CDF** (also used in (Lu, Angelopoulos, and Pomerantz 2022) as a baseline): A simple method to construct prediction intervals directly from cumulative probabilities.
- **Weighted CRC** (Xu, Guo, and Wei 2023): A conformal risk-control method for ordinal classification that weights miscoverage by class-dependent severity to produce calibrated interval predictions.

Evaluation Metrics. The evaluation metrics used in our experiments include the coverage of the true label and the average size of prediction sets under coverage, which are detailed as follows:

- **Coverage** (examine if $\geq 1 - \alpha$) is defined as the empirical coverage rate and can be computed by:

$$\text{Coverage} = \frac{1}{n} \sum_{i=n+1}^{n+m} \mathbb{1}[Y_i \in \hat{C}(X_i)],$$

where $D_{\text{test}} = \{(X_i, Y_i)\}_{i=n+1}^{n+m}$ is the test dataset with m being the number of test samples. \hat{C} denotes any ordinal CP methods (i.e., Naive CDF, Ordinal APS, min-CPS or min-RCPS) used in the experiment.

- **The average size of prediction sets (smaller better)** is defined as average number of classes in prediction sets:

$$\text{Prediction set size} = \frac{1}{m} \sum_{i=n+1}^{n+m} \sum_{y \in \mathcal{Y}} \mathbb{1}[y \in \hat{C}(X_i)],$$

α	Method	Coverage	Prediction Set Size
$\alpha = 0.1$	Naive CDF	0.9012 ± 0.0046	23.2533 ± 0.0693
	Ordinal APS	0.9024 ± 0.0030	15.1795 ± 0.0596
	WCRC	0.9014 ± 0.0043	15.1587 ± 4.9565
	min-CPS	0.9034 ± 0.0057	9.1179 ± 0.0964
	min-RCPS	0.9033 ± 0.0054	8.9235 ± 0.0510
$\alpha = 0.05$	Naive CDF	0.9507 ± 0.0042	25.5960 ± 0.1128
	Ordinal APS	0.9509 ± 0.0026	16.7082 ± 0.0648
	WCRC	0.9520 ± 0.0000	16.6785 ± 0.0523
	min-CPS	0.9517 ± 0.0036	10.6557 ± 0.0959
	min-RCPS	0.9516 ± 0.0036	10.7631 ± 0.0733
$\alpha = 0.01$	Naive CDF	0.9904 ± 0.0019	29.2091 ± 0.1967
	Ordinal APS	0.9903 ± 0.0018	19.6791 ± 0.1949
	WCRC	0.9905 ± 0.0000	19.7776 ± 0.0216
	min-CPS	0.9901 ± 0.0016	14.7986 ± 0.2855
	min-RCPS	0.9894 ± 0.0019	14.4887 ± 0.3399

Table 2: Coverage and prediction set size on Avocado Price with various α . (1) When $\alpha = 0.1$, the prediction set size decreases by 39.93%↓ for min-CPS and 41.21%↓ for min-RCPS. (2) When $\alpha = 0.05$, min-CPS and min-RCPS reduce the prediction set size by 36.22%↓ and 35.58%↓, respectively. (3) When $\alpha = 0.01$, the reduction is 24.80%↓ for min-CPS and 26.38%↓ for min-RCPS. On average across all three α values, min-CPS and min-RCPS reduce the prediction set size by 33.65%↓ and 34.39%↓, respectively.

which iterates all test samples and counts how many ordinal labels are included in the $\hat{C}(X_i)$. Smaller sets imply higher prediction certainty and informativeness.

Parameter Setting. We perform experiments under different miscoverage rates α (i.e., 0.1, 0.05, and 0.01) to thoroughly evaluate performance and robustness. For min-RCPS, we tune the length regularization hyper-parameter λ from a reasonable range $\{0, 0.001, 0.003, 0.005, \dots, 0.019\}$ (with 0.002 increasing step from 0.001), where we also conduct a sensitivity experiment later.

4.2 Experiment Results

Comparison with Baselines Across Benchmark Datasets.

The overall performance comparison of all ordinal CP methods across four benchmark datasets is reported in Table 1, where we follow a standard practice in the CP literature (Romano, Sesia, and Candes 2020; Angelopoulos et al. 2020; Huang et al. 2024) and set $\alpha = 0.1$. The results clearly demonstrate that both min-CPS and min-RCPS significantly improve the predictive efficiency over the best one of Ordinal APS, Naive CDF and WCRC.

Specifically, on Temperature, UTKFace, Avocado Price and IMDB, min-CPS relatively reduces the prediction set size by 5.11%↓, 7%↓, 39.93%↓ and 3.49%↓, while min-RCPS relatively reduces the prediction set size by 6.96%↓, 7.01%↓, 41.21%↓ and 3.16%↓. On average across all four datasets, min-CPS and min-RCPS reduce the prediction set size by 14%↓ and 15%↓, respectively. These relative improvements of predictive efficiency (more than 10% ↓ in prediction set size) are clearly significant over the CP literature (Ding et al. 2024; Huang et al. 2024; Shi et al. 2024).




Figure 1: Impact of the regularization hyper-parameter λ in min-RCPS on coverage and prediction set size on the Avocado Price dataset. When the coverage is guaranteed, the prediction set size initially decreases as λ increases from 0 to 0.003, where min-RCPS outperforms min-CPS by 2.13%↓ in terms of prediction set size. We find that setting λ to a relatively small value is sufficient to improve the predictive efficiency, and this happens on other datasets.




Figure 2: The empirical coverage rate $F(\tau)$: verify $F(\tau)$ monotonically increases in τ (ref. Section 3.3)

Sensitivity Analysis under Different Miscoverage Levels.

We further evaluate the methods under varying miscoverage levels, i.e., $\alpha \in \{0.1, 0.05, 0.01\}$, on the UTKFace dataset. We report the results of coverage and prediction set size for all methods in Table 2. Our results indicate that our methods consistently achieves smaller prediction sets across all evaluated α values, highlighting its stability and robustness.

Specifically, our observation can be summarized by

- When $\alpha = 0.1$, the prediction set size decreases by 39.93%↓ for min-CPS and 41.21%↓ for min-RCPS.
- When $\alpha = 0.05$, min-CPS and min-RCPS reduce the prediction set size by 36.22%↓ and 35.58%↓, respectively, compared with the best baseline.
- When $\alpha = 0.01$, the reduction is 24.80%↓ for min-CPS and 26.38%↓ for min-RCPS. On average across all three α values, min-CPS and min-RCPS reduce the prediction set size by 33.65%↓ and 34.39%↓, respectively.

Impact of Regularization Parameter λ in min-RCPS.

We analyze the effect of the regularization hyperparameter λ for min-RCPS on coverage and prediction set sizes. Figures 1 illustrate these effects on the Avocado Price datasets:

- As λ increases, the prediction set sizes initially decrease and reach a smallest prediction set size when $\lambda = 0.003$.
- Appropriately choosing λ significantly enhances the compactness of prediction sets as maintaining valid coverage.

Dataset	Method	Running Time (s)
Temperature	Ordinal APS	6183
	min-CPS	312
	min-RCPS	321
UTKFace	Ordinal APS	4855
	min-CPS	223
	min-RCPS	215
Avocado Price	Ordinal APS	711
	min-CPS	67
	min-RCPS	63
IMDB	Ordinal APS	823
	min-CPS	84
	min-RCPS	87

Table 3: Running time across datasets ($\alpha = 0.1$). All methods guarantee at least $1 - \alpha$ coverage rate with 10 trials. On Temperature, UTKFace, Avocado Price and IMDB, min-CPS speeds up x19.82, x21.77, x10.62 and x9.80 over Ordinal APS, while min-RCPS speeds up x19.26, x22.58, x11.29 and x9.46 over Ordinal APS. The stop criteria for both min-CPS and min-RCPS are to keep iterating until the coverage has been guaranteed (exceeding $1 - \alpha$).

- We recommend selecting the optimal λ via cross-validation or a calibration set for practical applications.

Running time comparison. We also compare the running time for all methods. Table 3 presents the running time comparison across four datasets under the setting $\alpha = 0.1$ with 10 trials. Our proposed methods, min-CPS and min-RCPS, demonstrate a significant computational advantage over the baseline Ordinal APS. Specifically, on Temperature, UTKFace, Avocado Price and IMDB, min-CPS speeds up x19.82, x21.77, x10.62 and x9.80 over Ordinal APS, while min-RCPS speeds up x19.26, x22.58, x11.29 and x9.46 over Ordinal APS. For the fair comparison, the stop criteria for both min-CPS and min-RCPS are to keep iterating until the coverage has been guaranteed (exceeding $1 - \alpha$).

5 Conclusion

In this work, we introduced **min-CPS**, a model-agnostic conformal prediction method for ordinal classification based on an instance-level minimum-length covering formulation, along with a linear-time sliding-window algorithm. Under a mild radial monotonicity condition, the resulting prediction sets are nested, enabling efficient calibration with valid marginal coverage. We further proposed **min-RCPS**, a length-regularized extension that incorporates interval-length information while preserving the same coverage guarantee. Across vision, sensor, and tabular datasets, both methods maintain target coverage while producing more compact prediction sets than existing ordinal CP approaches, reducing average interval sizes by **14%–15%** (up to **30–40%**) and achieving **10×–22×** speedups over Ordinal APS.

?

Acknowledgements

The authors gratefully acknowledge the in part support by the USDA-NIFA funded AgAID Institute award 2021-67021-35344, and the NSF grant CNS-2312125, IIS-2443828, DUE-2519063. The views expressed are those of the authors and do not reflect the official policy or position of the USDA-NIFA and NSF.

References

- Albuquerque, T.; Cruz, R.; and Cardoso, J. S. 2021. Ordinal losses for classification of cervical cancer risk. *PeerJ Computer Science*, 7: e457.
- Angelopoulos, A.; Bates, S.; Malik, J.; and Jordan, M. I. 2020. Uncertainty sets for image classifiers using conformal prediction. *arXiv preprint arXiv:2009.14193*.
- Angelopoulos, A. N.; and Bates, S. 2021. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. *arXiv preprint arXiv:2107.07511*.
- Barber, R. F.; Candès, E. J.; Ramdas, A.; and Tibshirani, R. J. 2021. Predictive inference with the jackknife+. *The Annals of Statistics*, 49(1): 486–507.
- Chakraborty, S.; Tyagi, C.; Qiao, H.; and Guo, W. 2024. Distribution-Free Conformal Prediction for Ordinal Classification. *arXiv preprint arXiv:2404.16610*.
- Chen, D.; Liu, Z.; Yang, C.; Wang, D.; Yan, Y.; Xu, Y.; and Ji, X. 2025. ConformalSAM: Unlocking the Potential of Foundational Segmentation Models in Semi-Supervised Semantic Segmentation with Conformal Prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 24045–24055.
- Cheng, Y.; Ying, H.; Hu, R.; Wang, J.; Zheng, W.; Zhang, X.; Chen, D.; and Wu, J. 2023. Robust image ordinal regression with controllable image generation. *arXiv preprint arXiv:2305.04213*.
- Cherian, J. J.; Gibbs, I.; and Candès, E. J. 2024. Large Language Model Validity via Enhanced Conformal Prediction Methods. *arXiv preprint arXiv:2406.09714*.
- Correia, A.; Massoli, F. V.; Louizos, C.; and Behboodi, A. 2024. An Information Theoretic Perspective on Conformal Prediction. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Cresswell, J. C.; Sui, Y.; Kumar, B.; and Vouitsis, N. 2024. Conformal Prediction Sets Improve Human Decision Making. In *International Conference on Machine Learning*, 9439–9457. PMLR.
- Dey, P.; Merugu, S.; and Kaveri, S. 2023. Conformal Prediction Sets for Ordinal Classification. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Ding, T.; Angelopoulos, A.; Bates, S.; Jordan, M.; and Tibshirani, R. J. 2024. Class-conditional conformal prediction with many classes. *Advances in Neural Information Processing Systems*, 36.
- Dong, Z.; Wu, Y.; Chen, C.; Zou, Y.; Zhang, Y.; and Zhou, J. H. 2025. Improve Representation for Imbalanced Regression through Geometric Constraints. *arXiv:2503.00876*.
- Doula, A.; Mühlhäuser, M.; and Guinea, A. S. 2024. AR-CP: Uncertainty-Aware Perception in Adverse Conditions with Conformal Prediction and Augmented Reality for Assisted Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 216–226.
- Feldman, S.; Bates, S.; and Romano, Y. 2021. Improving conditional coverage via orthogonal quantile regression. *Advances in neural information processing systems*, 34: 2060–2071.
- Fernandez-Navarro, F.; Campoy-Munoz, P.; Hervás-Martínez, C.; Yao, X.; et al. 2013. Addressing the EU sovereign ratings using an ordinal regression approach. *IEEE transactions on cybernetics*, 43(6): 2228–2240.
- Gendler, A.; Weng, T.-W.; Daniel, L.; and Romano, Y. 2021. Adversarially robust conformal prediction. In *International Conference on Learning Representations*.
- Ghosh, S.; Shi, Y.; Belkhouja, T.; Yan, Y.; Doppa, J.; and Jones, B. 2023. Probabilistically robust conformal prediction. In *UAI*, 681–690. PMLR.
- Gibbs, I.; Cherian, J. J.; and Candès, E. J. 2023. Conformal Prediction With Conditional Guarantees. *arXiv preprint arXiv:2305.12616*.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330. PMLR.
- Gutiérrez, P. A.; Perez-Ortiz, M.; Sanchez-Monedero, J.; Fernandez-Navarro, F.; and Hervás-Martínez, C. 2015. Ordinal regression methods: survey and experimental study. *IEEE Transactions on Knowledge and Data Engineering*, 28(1): 127–146.
- He, S.; Zhang, Y.; Xie, R.; Jiang, D.; and Ming, A. 2022. Rethinking Image Aesthetics Assessment: Models, Datasets and Benchmarks. In *IJCAI*, 942–948.
- Hirk, R.; Hornik, K.; and Vana, L. 2019. Multivariate ordinal regression models: an analysis of corporate credit ratings. *Statistical Methods & Applications*, 28(3): 507–539.
- Huang, J.; Cai, X.; Liu, K.; Cao, Y.; Wei, H.; and An, B. 2025. Conformal Prediction for Deep Classifier via Truncating.
- Huang, J.; Xi, H.; Zhang, L.; Yao, H.; Qiu, Y.; and Wei, H. 2024. Conformal Prediction for Deep Classifier via Label Ranking. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*.
- Kang, J. 2019. UTKFace Dataset (Kaggle version). <https://www.kaggle.com/datasets/jangedoo/utkface-new/data>. Accessed: 2025-07-21.
- Keramati, M.; Meng, L.; and Evans, R. D. 2024. ConR: Contrastive Regularizer for Deep Imbalanced Regression. *arXiv:2309.06651*.
- Kiggins, J. 2018. Avocado Prices. <https://www.kaggle.com/datasets/neuromusic/avocado-prices>. Accessed: 2025-07-21.
- Kirchgessner, W. 2018. Electric Motor Temperature. <https://www.kaggle.com/datasets/wkigrsn/electric-motor-temperature>. Accessed: 2025-07-21.

- Kiyani, S.; Pappas, G. J.; and Hassani, H. 2024. Length Optimization in Conformal Prediction. In *Annual Conference on Neural Information Processing Systems*.
- Kong, S.; Shen, X.; Lin, Z.; Mech, R.; and Fowlkes, C. 2016. Photo aesthetics ranking network with attributes and content adaptation. In *European conference on computer vision*, 662–679. Springer.
- Kwon, Y. S.; Han, I.; and Lee, K. C. 1997. Ordinal pairwise partitioning (OPP) approach to neural networks training in bond rating. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 6(1): 23–40.
- Lei, J.; G’Sell, M.; Rinaldo, A.; Tibshirani, R. J.; and Wasserman, L. 2018. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523): 1094–1111.
- Liu, Z.; Yufei, C.; Yan, Y.; Xu, Y.; Ji, X.; Liu, X.; and Chan, A. B. 2024. The Pitfalls and Promise of Conformal Inference Under Adversarial Attacks. In *Forty-first International Conference on Machine Learning*.
- Lu, C.; Angelopoulos, A. N.; and Pomerantz, S. 2022. Improving Trustworthiness of AI Disease Severity Rating in Medical Imaging with Ordinal Conformal Prediction Sets. *arXiv preprint arXiv:2207.02238*.
- Niu, Z.; Zhou, M.; Wang, L.; Gao, X.; and Hua, G. 2016. Ordinal regression with multiple output CNN for age estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4920–4928.
- Oliveira, R. I.; Orenstein, P.; Ramos, T.; and Romano, J. V. 2024. Split conformal prediction and non-exchangeable data. *JMLR*, 25(225): 1–38.
- Pedregosa, F.; Bach, F.; and Gramfort, A. 2017. On the consistency of ordinal regression methods. *Journal of Machine Learning Research*, 18(55): 1–35.
- Plassier, V.; Fishkov, A.; Guizani, M.; Panov, M.; and Moulines, E. 2024. Probabilistic Conformal Prediction with Approximate Conditional Validity. *arXiv preprint arXiv:2407.01794*.
- Romano, Y.; Patterson, E.; and Candes, E. 2019. Conformalized quantile regression. *Advances in neural information processing systems*, 32.
- Romano, Y.; Sesia, M.; and Candes, E. 2020. Classification with valid and adaptive coverage. *Advances in neural information processing systems*, 33: 3581–3591.
- Rothe, R.; Timofte, R.; and Gool, L. V. 2015. DEX: Deep EXpectation of apparent age from a single image. In *IEEE International Conference on Computer Vision Workshops (ICCVW)*.
- Shafer, G.; and Vovk, V. 2008. A Tutorial on Conformal Prediction. *Journal of Machine Learning Research*, 9(3).
- Shao, S.; Pei, Z.; Wu, X.; Liu, Z.; Chen, W.; and Li, Z. 2023. Iebins: Iterative elastic bins for monocular depth estimation. *Advances in Neural Information Processing Systems*, 36: 53025–53037.
- Shi, Y.; GHOSH, S.; Belkhouja, T.; Doppa, J.; and Yan, Y. 2024. Conformal Prediction for Class-wise Coverage via Augmented Label Rank Calibration. In *Annual Conference on Neural Information Processing Systems*.
- Shi, Y.; Shahrokhi, H.; Jia, X.; Chen, X.; Doppa, J.; and Yan, Y. 2025. Direct Prediction Set Minimization via Bilevel Conformal Classifier Training. In *Forty-second International Conference on Machine Learning*.
- Shin, N.; Lee, S.; and Kim, C. 2022. Moving Window Regression: A Novel Approach to Ordinal Regression. In *Proc. IEEE/CVF CVPR*, 18760–18769.
- Straitouri, E.; Wang, L.; Okati, N.; and Rodriguez, M. G. 2023. Improving expert predictions with conformal prediction. In *International Conference on Machine Learning*, 32633–32653. PMLR.
- Stutz, D.; Dvijotham, K. D.; Cemgil, A. T.; and Doucet, A. 2022. Learning Optimal Conformal Classifiers. In *International Conference on Learning Representations*.
- Su, J.; Luo, J.; Wang, H.; and Cheng, L. 2024. API Is Enough: Conformal Prediction for Large Language Models Without Logit-Access. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 979–995.
- van der Laan, L.; and Alaa, A. M. 2024. Self-Calibrating Conformal Prediction. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Vovk, V. 2012. Conditional validity of inductive conformal predictors. In *ACML*, 475–490. PMLR.
- Vovk, V. 2015. Cross-conformal predictors. *Annals of Mathematics and Artificial Intelligence*, 74(1): 9–28.
- Vovk, V.; Gammerman, A.; and Saunders, C. 1999. Machine-learning applications of algorithmic randomness.
- Vovk, V.; Gammerman, A.; and Shafer, G. 2005. *Algorithmic learning in a random world*, volume 29. Springer.
- Wang, J.; Chen, J.; Liu, J.; Tang, D.; Chen, D. Z.; and Wu, J. 2025. A Survey on Ordinal Regression: Applications, Advances, and Prospects. *arXiv preprint arXiv:2503.00952*.
- Wang, J.; Cheng, Y.; Chen, J.; Chen, T.; Chen, D. Z.; and Wu, J. 2023. Ord2Seq: Regarding Ordinal Regression as Label Sequence Prediction. In *Proc. IEEE/CVF ICCV*, 5865–5875.
- Xu, Y.; Guo, W.; and Wei, Z. 2023. Conformal Risk Control for Ordinal Classification. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2346–2355.
- Yu, Q.; Xie, J.; Nguyen, A.; Zhao, H.; Zhang, J.; Fu, H.; Zhao, Y.; Zheng, Y.; and Meng, Y. 2024. CLIP-DR: Textual knowledge-guided diabetic retinopathy grading with ranking-aware prompting. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 667–677. Springer.
- Zhang, M.; Hu, X.; Gu, L.; Liu, L.; Kobayashi, K.; Harada, T.; Yan, Y.; Summers, R. M.; and Zhu, Y. 2024. A New Benchmark: Clinical Uncertainty and Severity Aware Labeled Chest X-Ray Images With Multi-Relationship Graph Learning. *IEEE Transactions on Medical Imaging*.
- Zhang, Z.; Bao, J.; Zhou, Z.; Cheng, L.; Luo, R.; et al. 2025. Residual Reweighted Conformal Prediction for Graph Neural Networks. In *The 41st Conference on Uncertainty in Artificial Intelligence*.