

Symmetric Aggregation of Conformity Scores for Efficient Uncertainty Sets

Nabil Alami^{1*}, Jad Zakharia^{2*}, Souhaib Ben Taieb^{1 †}

¹Department of Statistics and Data Science, Mohamed Bin Zayed University of Artificial Intelligence (MBZUAI)

²Ecole Normale Supérieure de Cachan

nabil.alami@mbzuai.ac.ae, jad.zakharia@enpc.fr, souhaib.bentaieb@mbzuai.ac.ae

Abstract

Access to multiple predictive models trained for the same task, whether in regression or classification, is increasingly common in many applications. Aggregating their predictive uncertainties to produce reliable and efficient uncertainty quantification is therefore a critical but still underexplored challenge, especially within the framework of conformal prediction (CP). While CP methods can generate individual prediction sets from each model, combining them into a single, more informative set remains a challenging problem. To address this, we propose SACP (Symmetric Aggregated Conformal Prediction), a novel method that aggregates nonconformity scores from multiple predictors. SACP transforms these scores into e-values and combines them using any symmetric aggregation function. This flexible design enables a robust, data-driven framework for selecting aggregation strategies that yield sharper prediction sets. We also provide theoretical insights that help justify the validity and performance of the SACP approach. Extensive experiments on diverse datasets show that SACP consistently improves efficiency and often outperforms state-of-the-art model aggregation baselines.

Code — <https://github.com/jadz1/SACP-SymmetricConformalAggregation>

1 Introduction

Artificial Intelligence (AI) has transformed numerous domains, including computer vision, natural language processing, forecasting, and decision-making in high-stakes environments. Its success largely arises from the ability to learn complex patterns and extract meaningful representations from large, high-dimensional datasets. However, despite their impressive predictive performance, many AI models fail to provide reliable estimates of uncertainty (Abdar et al. 2021; Wang et al. 2025). This shortcoming is particularly critical in high-risk applications, where erroneous predictions can lead to severe consequences. In such settings, it is essential to assess not only what a model predicts but also how confident it is in those predictions.

*These authors contributed equally.

†Department of Computer Science, University of Mons
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In recent years, conformal prediction (CP) has emerged as a powerful, distribution-free framework for constructing prediction sets with finite-sample coverage guarantees (Vovk, Gammerman, and Shafer 2005; Angelopoulos and Bates 2022). After training a model on the training data, CP uses a separate calibration set to compute nonconformity scores (NCSs), which measure how well each sample conforms to, or deviates from, the patterns learned by the model. These calibration scores are assumed to be exchangeable with the test data. Under this assumption, CP provides exact marginal coverage guarantees without relying on strong distributional assumptions, making it particularly valuable in applications where reliability is critical.

When multiple predictive models are available for the same task, aggregating their outputs is a common strategy in machine learning to improve predictive performance, enhance robustness, and reduce variance (Dietterich 2000; Breiman 1996). However, despite the success of model aggregation in traditional settings, its integration within the CP framework remains relatively underexplored. A central challenge lies in effectively combining the outputs of multiple conformal predictors to produce sharper prediction sets while preserving the exact coverage guarantees that make CP so appealing. In this context, efficiency refers to the tightness or expected size of the prediction sets, where smaller sets at the same coverage level indicate a more informative and useful predictor. This motivates two key questions: *How can conformal predictions from multiple models be aggregated? And does such aggregation improve coverage and/or efficiency?*

Several methods have been proposed to address this challenge. One prominent line of work focuses on model selection, aiming to identify the most suitable predictors based on specific performance criteria (Jin and Candès 2023; Yang and Kuchibhotla 2024; Gasparin and Ramdas 2024a). Other studies have explored merging prediction sets from individual conformal predictors in a black-box manner (Gasparin and Ramdas 2024b). More recent approaches go further by leveraging the structure of the conformity scores produced by each model (Rivera, Patel, and Tewari 2025; Luo and Zhou 2025). These methods aim to exploit shared patterns across predictors to improve efficiency, producing sharper prediction sets while maintaining valid coverage guarantees. However, some approaches rely on additional hyperparam-

ters or fail to fully utilize all available data. Furthermore, the existing literature lacks a systematic comparison of these aggregation strategies, leaving their relative strengths and limitations insufficiently understood.

To address these gaps, we propose Symmetric Aggregated Conformal Prediction (SACP), a novel method that constructs a single prediction set by symmetrically aggregating normalized conformity scores from multiple predictors. SACP operates in two stages. First, it normalizes the scores using a transformation inspired by recent advances in e-value-based CP (Balinsky and Balinsky 2024; Gauthier, Bach, and Jordan 2025b,a). Second, it applies an arbitrary symmetric aggregation function to combine the normalized scores. This design yields a flexible framework that supports data-driven aggregation strategies capable of preserving coverage while enhancing efficiency. The main contributions of this work are summarized as follows:

- We introduce **Symmetric Aggregated Conformal Prediction (SACP)**, a novel method that symmetrically combines normalized nonconformity scores from multiple predictors to construct a single, informative prediction set.
- We develop a data-driven variant of SACP, supported by theoretical analysis, that adaptively selects aggregation strategies to improve efficiency while preserving exact marginal coverage.
- We conduct a comprehensive empirical evaluation on both regression and classification tasks, demonstrating that SACP consistently produces sharper prediction sets and outperforms existing conformal aggregation baselines.

2 Background

We consider a supervised learning setting with a dataset $\{(X_i, Y_i)\}_{i \in \mathcal{I}} \subset \mathcal{X} \times \mathcal{Y}$ drawn i.i.d. from an unknown distribution P , where \mathcal{X} is the input space, $\mathcal{Y} \subset \mathbb{R}$ is the output space, and \mathcal{I} is the set of data indices. We split the index set \mathcal{I} into three disjoint subsets: $\mathcal{I}_{\text{train}}$, \mathcal{I}_{cal} , and $\mathcal{I}_{\text{test}}$, corresponding respectively to the training, calibration, and test data. We denote $\mathcal{I}_{\text{cal}} = \{1, \dots, n\}$, where $n = |\mathcal{I}_{\text{cal}}|$ is the calibration set size. We consider K base regression or classification predictors $\hat{\mu}^{(1)}, \dots, \hat{\mu}^{(K)}$, each trained on $\{(X_i, Y_i)\}_{i \in \mathcal{I}_{\text{train}}}$. Given a new test pair $(X_{\text{test}}, Y_{\text{test}})$, assumed exchangeable with the calibration set $\{(X_i, Y_i)\}_{1 \leq i \leq n}$, our goal is to construct a prediction set for the unknown label $Y_{\text{test}} \in \mathcal{Y}$ with a guaranteed miscoverage level $\alpha \in (0, 1)$ while minimizing its expected length.

Conformal Prediction. Conformal Prediction provides distribution-free prediction sets with finite-sample coverage guarantees (Vovk, Gammerman, and Shafer 2005; Angelopoulos and Bates 2022). In our setting, for each predictor $k = 1, \dots, K$, we compute *nonconformity scores* (NCSs) on the calibration set $\{(X_i, Y_i)\}_{1 \leq i \leq n}$, defined as $s_i^{(k)} = s^{(k)}(X_i, Y_i)$, where $s^{(k)} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ measures how atypical the true label Y_i is with respect to the model prediction $\hat{\mu}^{(k)}(X_i)$. For regression, we set $s^{(k)}(X, Y) = |Y - \hat{\mu}^{(k)}(X)|$, and for classification, $s^{(k)}(X, Y) = 1 -$

$\hat{\mu}_Y^{(k)}(X)$, where $\hat{\mu}^{(k)}(X)$ and $\hat{\mu}_Y^{(k)}(X)$ denote the estimated regression function and the predicted probability of class Y , respectively. Alternative definitions of nonconformity scores can also be used; see Dheur et al. (2025) for examples. Denote by $s_{(1)}^{(k)} \leq \dots \leq s_{(n)}^{(k)}$ the order statistics of the calibration scores, and let $\hat{Q}_\alpha^k = s_{(\lceil (1-\alpha)(n+1) \rceil)}^{(k)}$ be their empirical quantile, computed from $\mathcal{S}_{\text{cal}} := \{s_i^{(k)}\}_{1 \leq i \leq n}$. The corresponding prediction set is defined as

$$\mathcal{C}_\alpha^k(X_{\text{test}}) = \{y \in \mathcal{Y} \mid s^{(k)}(X_{\text{test}}, y) \leq \hat{Q}_\alpha^k\}, \quad (1)$$

which guarantees finite-sample marginal coverage:

$$\mathbb{P}(Y_{\text{test}} \in \mathcal{C}_\alpha^k(X_{\text{test}})) \geq 1 - \alpha. \quad (2)$$

Alternatively, one can define $\tilde{s}^{(k)} = -s^{(k)}$ and express the prediction set as $\mathcal{C}_\alpha^k(X_{\text{test}}) = \{y \in \mathcal{Y} \mid \tilde{s}^{(k)}(X_{\text{test}}, y) \geq q_\alpha^k\}$, with $q_\alpha^k = \tilde{s}_{(\lfloor (n+1)\alpha \rfloor)}^{(k)}$, which still satisfies (2). This formulation highlights that the resulting prediction set depends on the choice of nonconformity score.

CP Aggregation Methods. Aggregating predictors within the CP framework aims to combine the outputs of multiple models trained for the same task into a single unified prediction set. Since efficiency is also a key objective, aggregation methods typically aim to produce the smallest possible prediction sets while maintaining valid coverage guarantees. Broadly, these approaches can be grouped into two main classes:

1. Combining prediction sets. A simple aggregation strategy combines the individual prediction sets either by *intersection*, $\bigcap_{k=1}^K \mathcal{C}_\alpha^k$, which reduces the set size but yields coverage of at least $1 - K\alpha$, or by *union*, $\bigcup_{k=1}^K \mathcal{C}_\alpha^k$, which guarantees coverage at the cost of a larger prediction set. Yang and Kuchibhotla (2024) proposed a model selection approach that chooses the single best predictor by minimizing the expected prediction set length:

$$\mathcal{C}_\alpha^{\text{Sel}}(X_{\text{test}}) := \mathcal{C}_\alpha^{k^*}(X_{\text{test}}), \quad k^* \in \arg \min_{1 \leq k \leq K} \mathbb{E}[|\mathcal{C}_\alpha^k(X_{\text{test}})|]. \quad (3)$$

A more intuitive alternative is the *majority vote* method introduced by Gasparin and Ramdas (2024b), which includes a candidate label if it is accepted by most individual predictors:

$$\mathcal{C}_\alpha^M(X_{\text{test}}) := \left\{ y \in \mathcal{Y} \mid \frac{1}{K} \sum_{k=1}^K \mathbf{1}_{[y \in \mathcal{C}_\alpha^k(X_{\text{test}})]} > \frac{1}{2} \right\}, \quad (4)$$

achieving coverage of at least $1 - 2\alpha$. To further improve efficiency, the authors proposed a randomized variant in which the fixed majority threshold $1/2$ is replaced by $1/2 + U$, with $U \sim \text{Unif}[0, 1]$. While elegant, the majority vote approach operates solely on the final prediction sets, limiting its ability to fully exploit the information available at the score level.

2. Combining scores. Another class of methods aggregates information directly from the nonconformity scores (NCSs). Luo and Zhou (2025) propose aggregating NCSs through a convex combination. For any pair (X, Y) , let

Method	Coverage Guarantee	Time complexity	Calibration	Notes
Wagg	$1 - \alpha$	$\mathcal{O}(W_{\text{grid}} \cdot n_1 + n_2 \cdot W_{\text{grid}} \cdot D_{\text{grid}} + n_3)$	Split in 2 or 3	Challenging in high dimensions
CSA	$1 - \alpha$	$\mathcal{O}(M \cdot K \cdot n)$	Split in 2	Based on projected quantiles
CM & CR	$1 - 2\alpha$	$\mathcal{O}(K \cdot n)$	No split	Set-level symmetric aggregation
SACP	$1 - \alpha$	$\mathcal{O}(D_{\text{grid}} \cdot K \cdot n)$	No split	Choice of aggregating function

Table 1: Comparison of aggregation methods. Time complexity refers to the time to generate the prediction set of a single test point. **Notation:** n_i — calibration size of subset i of the calibration set; n — total calibration size; K — number of predictors; W_{grid} — size of the grid for weights; M — number of directions. D_{grid} — number of classes (in classification) or length of the discretized target space \mathcal{Y} (in regression).

$s(X, Y) = (s^{(1)}(X, Y), \dots, s^{(K)}(X, Y))^\top$ denote the corresponding vector of NCSs. The aggregated score is then given by $s(X, Y)^\top \mathbf{w}$, where \mathbf{w} lies on the $(K - 1)$ -simplex (i.e., $w_k \geq 0$ and $\sum_{k=1}^K w_k = 1$) and is chosen to minimize the expected prediction set length. The corresponding prediction set is defined as

$$\mathcal{C}_\alpha^{\mathbf{w}}(X_{\text{test}}) := \left\{ y \in \mathcal{Y} \mid s(X_{\text{test}}, y)^\top \mathbf{w} \leq \hat{Q}_\alpha \right\}, \quad (5)$$

where \hat{Q}_α is the empirical $(1 - \alpha)$ -quantile of the aggregated scores, computed from a subset of the calibration data to ensure valid coverage.

Rivera, Patel, and Tewari (2025) propose *Conformal Score Aggregation (CSA)*, a method based on multivariate quantiles. The calibration set is split into two subsets, $\mathcal{S}_{\text{cal}} := \mathcal{S}^{(1)} \cup \mathcal{S}^{(2)}$. On $\mathcal{S}^{(1)}$, CSA samples M directions $\{u_m\}_{m=1}^M$ on the positive unit sphere and determines thresholds q_m such that a fraction β of projections $u_m^\top s$ exceed each q_m . These thresholds define the envelope

$$H(\beta) = \bigcap_{m=1}^M \left\{ s \in \mathcal{S}^{(1)} : u_m^\top s \leq q_m(\beta) \right\},$$

and a binary search over β is performed to achieve $1 - \alpha$ coverage on $\mathcal{S}^{(1)}$. Next, each score $s \in \mathcal{S}^{(2)}$ is mapped to $\max_m \left(\frac{u_m^\top s}{q_m(\beta^*)} \right)$; rescaling $H(\beta^*)$ by the empirical $(1 - \alpha)$ -quantile of these normalized values preserves exchangeability and guarantees exact coverage. Table 1 summarizes these aggregation methods.

Setup Justification. In the context of aggregation, the key components are the training data, the calibration data, and the nonconformity scores (NCSs). Differences in any of these elements can lead to distinct CP settings. To ensure a fair and consistent comparison across methods, our setup uses K predictors trained on the same training dataset, evaluated on the same calibration set using a fixed nonconformity score.

3 Our SACP Method

Our SACP method aims to construct a single final prediction set from K base predictors. Recent advances in conformal prediction have underscored the importance of e -values for uncertainty quantification (Gauthier, Bach, and Jordan 2025b). Building on this idea, SACP normalize the nonconformity scores by transforming them into e -variables. They

are then combined through a symmetric aggregation function, enabling the use of standard (split) conformal prediction.

Definition 3.1. An e -variable E is a nonnegative random variable ($E \geq 0$) that satisfies, under a null hypothesis \mathbb{H}_0 ,

$$\mathbb{E}_{\mathbb{H}_0}[E] \leq 1. \quad (6)$$

An e -value is the observed value of an e -variable.

The first step of SACP is to construct e -variable-like transformations of the nonconformity scores, inspired by the work of Balinsky and Balinsky (2024), although we do not rely on the general e -value theory. This construction is based on the following proposition.

Proposition 3.2. Consider exchangeable and positive random variables $\{S_i\}_{1 \leq i \leq n+1}$. Then, the random variables

$$E_j = \frac{S_j}{\frac{1}{n+1} \sum_{i=1}^{n+1} S_i}, \quad j = 1, \dots, n \quad (7)$$

have expectation equal to one.

Specifically, for each predictor and its corresponding scores, SACP constructs $n + 1$ exchangeable e -variables of the form (7).

Let $y \in \mathcal{Y}$ denote a candidate label associated with X_{test} , so that $s^{(k)}(X_{\text{test}}, y)$ quantifies its conformity under predictor k . For each $k = 1, \dots, K$ and each calibration point (X_i, Y_i) , $i = 1, \dots, n$, we construct the calibration e -variable

$$E_i^{(k)}(y) = \frac{s^{(k)}(X_i, Y_i)}{\frac{1}{n+1} \left(\sum_{j=1}^n s^{(k)}(X_j, Y_j) + s^{(k)}(X_{\text{test}}, y) \right)}. \quad (8)$$

This quantity is the ratio of the i^{th} score from predictor k to the average of all calibration scores augmented with the test score for the candidate y .

Since the scores $\{s^{(k)}(X_i, Y_i)\}_{1 \leq i \leq n}$ are exchangeable, $E_i^{(k)}$ follows the same form as in (7) and is therefore an e -variable (under the null hypothesis $\mathbb{H}_0 : y = Y_{\text{test}}$, and assuming the NCS are positive). For each calibration point (X_i, Y_i) , we define its corresponding e -vector as $\mathbf{E}_i = (E_i^{(1)}, \dots, E_i^{(K)})^\top \in \mathbb{R}^K$. Similarly, the test e -variable is defined as

$$E_{\text{test}}^{(k)}(y) = \frac{s^{(k)}(X_{\text{test}}, y)}{\frac{1}{n+1} \left(\sum_{i=1}^n s^{(k)}(X_i, Y_i) + s^{(k)}(X_{\text{test}}, y) \right)}, \quad (9)$$

and the corresponding test e -vector is $\mathbf{E}_{\text{test}} = (E_{\text{test}}^{(1)}, \dots, E_{\text{test}}^{(K)})^\top$.

Note that all $\{E_i^{(k)}\}_{1 \leq i \leq n}$ share the same denominator, so their relative ordering is preserved. Moreover, the test e -variable mirrors the behavior of the test score, decreasing as $s^{(k)}(X_{\text{test}}, y)$ decreases.

The second step of SACP is to combine the constructed e -variables using a *symmetric* function $f : \mathbb{R}^K \rightarrow \mathbb{R}$, which merges the e -variables into new aggregated scores. For $i = 1, \dots, n$, we define

$$F_i(y) := f(\mathbf{E}_i(y)), \quad F_{\text{test}}(y) := f(\mathbf{E}_{\text{test}}(y)). \quad (10)$$

Working with symmetric aggregators is especially natural in a distribution-free setting and offers two key advantages. First, the aggregated score is invariant to how the base models are indexed or labeled: permuting their indices leaves the aggregated scores unchanged, ensuring that model labeling has no effect. Second, symmetry confines us to a structured and interpretable class of functions, which in turn makes it easier (a) to optimize the aggregator in practice and (b) to derive theoretical results on the resulting prediction-set size.

Note that the aggregated scores' behavior directly depends on the choice of the aggregation function. When the aggregated scores exhibit the characteristics of a classical NCS, in which smaller values correspond to higher conformity, the prediction set is defined as

$$C_\alpha^{f, \uparrow}(X_{\text{test}}) = \{y \in \mathcal{Y} \mid F_{\text{test}}(y) \leq \hat{Q}_\alpha(y)\}, \quad (11)$$

where

$$\hat{Q}_\alpha(y) := F_{(\lceil (1-\alpha)(n+1) \rceil)}(y) \quad (12)$$

is the upper empirical quantile of $\{F_i\}_{1 \leq i \leq n}$. In contrast, when higher values indicate better conformity, the prediction set is

$$C_\alpha^{f, \downarrow}(X_{\text{test}}) = \{y \in \mathcal{Y} \mid F_{\text{test}}(y) \geq \hat{q}_\alpha(y)\}, \quad (13)$$

where

$$\hat{q}_\alpha(y) := F_{(\lfloor \alpha(n+1) \rfloor)}(y) \quad (14)$$

is the lower empirical quantile of $\{F_i\}_{1 \leq i \leq n}$.

Theorem 3.3. *The SACP prediction sets defined in (11) and (13) verify the coverage property (2).*

The proof follows from the exchangeability of $\{F_i\}_{1 \leq i \leq n} \cup F_{\text{test}}$, see Appendix A. In our implementation of SACP, the default choice for f is the sum of its inputs. An overview of the main steps of SACP is presented in Figure 1.

Remark 1. If we interpret the original NCS as “smaller values correspond to better conformity,” then indeed the behavior of the aggregated score aligns with the monotonicity of f . If f is non-decreasing component-wise, the resulting merged score behaves as a NCS: lower values indicate better conformity. In this case, the prediction set is constructed using (11). In contrast, if f is decreasing component-wise, the interpretation is reversed: higher values indicate better conformity, and the prediction set is built according to (13). If f is non-monotonic, it is difficult to draw precise conclusions about the behavior of the aggregated scores.

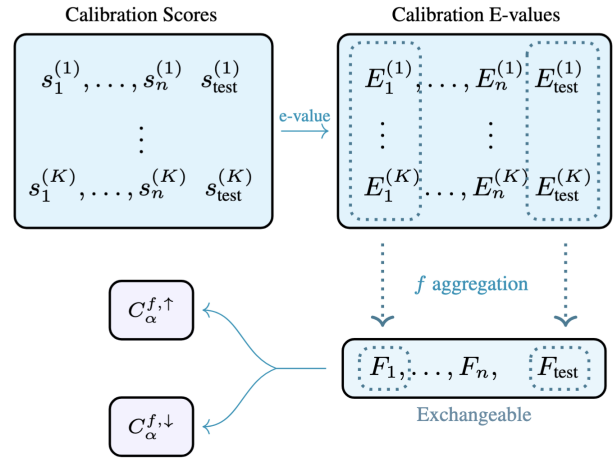


Figure 1: Diagram illustrating key steps of SACP

Remark 2. Unlike the standard CP framework, the thresholds \hat{q}_α and \hat{Q}_α depend on the candidate label y . In classification tasks, a separate quantile must be computed for each class, whereas in regression we discretize the output space \mathcal{Y} into a uniform grid and compute one quantile for each grid point.

Our SACP method offers several key advantages and is particularly relevant in the context of conformal aggregation:

- To the best of our knowledge, this is the first approach to apply symmetric conformal aggregation at the score level while guaranteeing $1 - \alpha$ coverage. This is achieved without requiring any additional data splitting of the calibration set (see Table 1).
- Transforming raw scores into e -variables standardizes them across models by ensuring a common first moment (i.e., an expected value equal to one), enabling fair aggregation regardless of differences in scale or distribution.
- The method is conceptually simple and allows a flexible choice of the symmetric function, paving the way for further improvements through data-driven aggregation strategies.

Theoretical analysis

The aggregation of scores using an arbitrary symmetric function is subtle and requires careful analysis. In this section, we examine the influence of the symmetric aggregation function and explore key properties of the empirical quantile. These results allow us to derive a worst-case bound for our prediction set in regression tasks. All proofs are provided in Appendix A.

The Aggregating Function. SACP requires selecting a symmetric function; however, not all symmetric functions yield more informative prediction sets. Since analyzing the general structure of symmetric functions is inherently difficult, we follow the framework of Zaheer et al. (2018) and

focus on the following subclass:

$$f(\mathbf{x}) := (\rho \circ \Phi)(\mathbf{x}) = \rho \left(\sum_{k=1}^K \phi(x_k) \right), \quad (15)$$

where $\mathbf{x} = (x_1, \dots, x_K)^T \in (\mathbb{R}_+)^K$, and ϕ and ρ are continuous scalar transformations.

As discussed earlier, the monotonicity of f plays a crucial role in determining how the aggregated scores behave like NCS. Therefore, we restrict attention to strictly monotonic functions ρ and ϕ .

If f is increasing, then ρ and ϕ share the same monotonicity. In this case, the aggregated scores F_i in (10) behave like NCS, and the corresponding prediction set is given in (11). Conversely, if f is decreasing, ρ and ϕ have opposite monotonicity, and the associated prediction set is given in (13). Monotonic transformations are particularly convenient, as their influence on the ordering of scores is simpler to analyze.

Proposition 3.4. *Let $V_{(1)} \leq \dots \leq V_{(n)}$ be the order statistics of $V_1, \dots, V_n \in \mathbb{R}$. Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be continuous and define $\tilde{V}_i = g(V_i)$; denote their order statistics by $\tilde{V}_{(1)} \leq \dots \leq \tilde{V}_{(n)}$. Then, for any $\alpha \in (0, 1)$,*

- If g is strictly decreasing, then

$$\tilde{V}_{(\lceil (n+1)(1-\alpha) \rceil)} = g(V_{(\lfloor (n+1)\alpha \rfloor)}). \quad (16)$$

- If g is strictly increasing, then

$$\tilde{V}_{(\lceil (n+1)(1-\alpha) \rceil)} = g(V_{(\lceil (n+1)(1-\alpha) \rceil)}). \quad (17)$$

This proposition directly implies the following result.

Proposition 3.5. *For any $\alpha \in (0, 1)$, consider a symmetric function of the form (15), where ρ and ϕ are monotonic. Then, for a test point X_{test} , we have:*

- If ρ is strictly non-decreasing:

$$\mathcal{C}_\alpha^{\rho \circ \Phi, \uparrow}(X_{test}) = \mathcal{C}_\alpha^{\Phi, \uparrow}(X_{test}), \quad (18)$$

$$\mathcal{C}_\alpha^{\rho \circ \Phi, \downarrow}(X_{test}) = \mathcal{C}_\alpha^{\Phi, \downarrow}(X_{test}). \quad (19)$$

- If ρ is strictly non-increasing:

$$\mathcal{C}_\alpha^{\rho \circ \Phi, \uparrow}(X_{test}) = \mathcal{C}_\alpha^{\Phi, \downarrow}(X_{test}), \quad (20)$$

$$\mathcal{C}_\alpha^{\rho \circ \Phi, \downarrow}(X_{test}) = \mathcal{C}_\alpha^{\Phi, \uparrow}(X_{test}). \quad (21)$$

Therefore, these results show that the dependence on ρ can be eliminated when constructing our prediction set: $\mathcal{C}_\alpha^{\rho \circ \Phi} = \mathcal{C}_\alpha^{\Phi}$.

For this reason, we consider the following class of symmetric functions.

Definition 3.6. *A function Φ is called a ϕ -aggregating function if it can be expressed as*

$$\Phi(\mathbf{x}) := \sum_{k=1}^K \phi(x_k), \quad (22)$$

for all $\mathbf{x} = (x_1, \dots, x_K)^T \in (\mathbb{R}_+)^K$, where $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}$ is continuous and monotonic. The set of all ϕ -aggregating functions is denoted by \mathcal{F}_{agg} .

From this point onward, we restrict our attention to aggregating functions of the form given in Definition 3.6. We next establish a bound on the length of our prediction set, which requires intermediate results on the empirical quantile.

Quantile Aggregation. Aggregating quantiles under arbitrary dependence is a classical problem in statistics and probability theory, with early work dating back to the early 20th century (Vincent 1912; Thomas and Ross 1980). More recent research has established theoretical bounds and relationships between individual quantiles and their aggregated counterparts (Lichtendahl Jr, Grushka-Cockayne, and Winkler 2013; Blanchet et al. 2024).

The SACP methods involves computing the empirical quantile of the aggregated scores, which we will seek to upper bound in order to establish our final theorem.

We now present an upper bound on the length of the SACP prediction set. Since the length naturally depends on the choice of the function ϕ in (22), quantifying this dependence exactly is often intractable. Instead, we derive a **worst-case bound** that holds uniformly over all admissible choices of ϕ . This bound characterizes the maximal width that any aggregated prediction set can attain, ensuring that any specific choice of ϕ can only improve (i.e., tighten) the prediction set relative to this limit.

Theorem 3.7 (Worst-case bound). *For a regression setting having $K > 1$ predictors, and using the absolute residual nonconformity scores, consider $\Phi \in \mathcal{F}_{agg}$. For a test point X_{test} , denote by $|\mathcal{C}_\alpha^\Phi(X_{test})|$ the length of the SACP prediction set, and by $|\mathcal{C}_\alpha^k(X_{test})|$ the length of the conformal prediction set of the k -th predictor $\hat{\mu}^{(k)}$ at level $\alpha \in [\frac{K}{n+1}, 1)$. Let $\alpha' := \alpha/K$, and define the model disagreement at X_{test} as*

$$\Delta_{test} = \max_{1 \leq k \leq K} \hat{\mu}^{(k)}(X_{test}) - \min_{1 \leq k \leq K} \hat{\mu}^{(k)}(X_{test}). \quad (23)$$

Then,

$$|\mathcal{C}_\alpha^\Phi(X_{test})| \leq \Delta_{test} + \max_k |\mathcal{C}_{\alpha'}^k(X_{test})|. \quad (24)$$

The proof of Theorem 3.7 relies on bounding the quantities that appear in the SACP prediction set. The inequality (24) provides a worst-case guarantee on the length of the aggregated prediction set.

An efficiency-oriented SACP method

We now propose an efficiency-oriented extension of SACP, building upon the theoretical guarantees established earlier. By Theorem 3.3, the prediction set \mathcal{C}_α^Φ achieves the desired $1 - \alpha$ coverage for any symmetric aggregation function $\Phi \in \mathcal{F}_{agg}$. This invariance allows us to exploit the flexibility in choosing Φ to improve efficiency, specifically, to minimize the expected length of the prediction sets.

Formally, we aim to find an aggregation function Φ^* that minimizes the prediction set length

$$\Phi^* = \arg \min_{\Phi \in \mathcal{F}_{agg}} |\mathcal{C}_\alpha^\Phi|. \quad (25)$$

However, searching over the entire space \mathcal{F}_{agg} is intractable. We therefore restrict attention to the parametric subclass

$$\Phi_p(\mathbf{x}) = \sum_{k=1}^K (x_k)^p, \quad p \in \mathbb{R}, \quad (26)$$

which encompasses several standard aggregation schemes. For instance, as $p \rightarrow +\infty$, Φ_p approaches the maximum operator, while $p \rightarrow -\infty$ yields the minimum. This family has also been successfully employed in conformal prediction for efficiency optimization (Braun et al. 2025).

We then select the exponent p^* that minimizes the average prediction set length on the (unlabeled) test set:

$$p^* = \arg \min_{p \in \mathbb{R}} \frac{1}{|\mathcal{I}_{\text{test}}|} \sum_{i \in \mathcal{I}_{\text{test}}} |C_{\alpha^p}^{\Phi_p}(X_i)|. \quad (27)$$

Since Theorem 3.3 guarantees valid coverage for any symmetric aggregator, this optimization preserves coverage while enhancing efficiency. Among all coverage-preserving functions Φ_p , we thus select the exponent p^* that minimizes the average set length. We refer to this enhanced variant as **SACP++**. Recall that the default SACP corresponds to $p = 1$.

4 Related Work

Several approaches have been proposed to aggregate conformal predictions. As discussed in Section 2, some methods focus on aggregating *prediction sets* (Yang and Kuchibhotla 2024; Gasparin and Ramdas 2024b), while others operate directly on the *scores* (Luo and Zhou 2025; Rivera, Patel, and Tewari 2025). We next review additional contributions in this area, highlighting their respective strengths and limitations.

A number of works in the literature define multivariate nonconformity scores for conformal prediction. Although not explicitly designed for aggregation, these methods naturally extend to it by mapping scores from \mathbb{R}^K to \mathbb{R} . A recent line of research leverages *Optimal Transport* to address multivariate score spaces (Klein et al. 2025; Thurin, Najahi, and Boyer 2025). These approaches construct transport maps that project multivariate nonconformity scores onto a single radial axis, enabling the application of standard univariate calibration procedures. While conceptually elegant and theoretically well-grounded, they can be computationally intensive and demand careful algorithmic design to ensure numerical stability and robust performance.

Our method involves the construction of nonconformity scores inspired by *e-values*. E-values have recently gained attention in the conformal prediction community as a powerful alternative to p-values (Vovk 2025; Wang and Ramdas 2022; Ramdas and Wang 2024; Balinsky and Balinsky 2024; Gauthier, Bach, and Jordan 2025b,a). Rooted in hypothesis testing, e-values enable more natural and admissible aggregation schemes (Vovk and Wang 2021; Wang 2024). They are particularly appealing due to their flexibility, anytime validity, robustness, post-hoc guarantees, and deep connections to betting and martingale theory. In contrast, their traditional counterparts, p-values, have long served as the cornerstone of conformal prediction but exhibit limitations when multiple p-values must be combined under arbitrary dependence structures (Vovk and Wang 2020; Vovk, Wang, and Wang 2022; Gasparin, Wang, and Ramdas 2025; Ramdas

and Wang 2024). While certain assumptions, such as exchangeability, can partially restore statistical power through refined inequalities (e.g., Gasparin and Ramdas 2025), p-values remain less flexible for aggregation in general settings.

5 Experiments and Results

We conduct a large-scale experimental study evaluating our methods, **SACP** and **SACP++**, on both regression and classification tasks, and compare their performance against several conformal aggregation techniques.

Datasets. For regression, we use a benchmark of single-output regression datasets from OpenML (Vanschoren et al. 2014), previously used in CP studies (Rivera, Patel, and Tewari 2025). For classification, we evaluate on **CIFAR-10** (Krizhevsky 2009) and **MNIST** (LeCun, Cortes, and Burges 2010), two standard benchmarks widely used in conformal prediction research (Luo and Zhou 2025, 2024).

Baselines. We compare against four aggregation methods described in Section 2: (i) the weighted aggregation approach of Luo and Zhou (2025), denoted **Wagg**; (ii) the multivariate-quantile method of Rivera, Patel, and Tewari (2025), denoted **CSA**; (iii) the deterministic and randomized majority-vote merging strategies of Gasparin and Ramdas (2024b), denoted **CM** and **CR**, respectively; and (iv) a best-model selection rule, denoted **BL**, which selects the individual model achieving the smallest average prediction-set length (Yang and Kuchibhotla 2024; Gasparin and Ramdas 2024a).

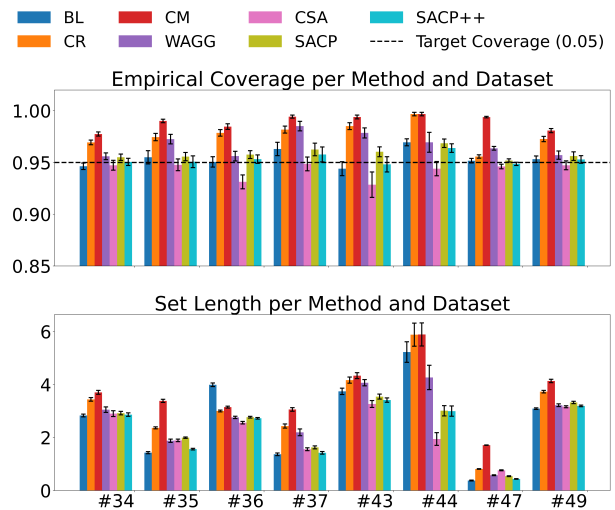


Figure 2: coverage and average prediction set length per method across OpenML regression datasets. For each dataset, the bars correspond, from left to right, to BL, CR, CM, Wagg, CSA, SACP, and SACP++.

Experimental setup. For each dataset, we perform 20 random splits into 80% training, 10% calibration, and 10% test

CIFAR-10					MNIST				
	Cov $\alpha=0.05$	Cov $\alpha=0.1$	Len $\alpha=0.05$	Len $\alpha=0.1$		Cov $\alpha=0.05$	Cov $\alpha=0.1$	Len $\alpha=0.05$	Len $\alpha=0.1$
<i>Base Models</i>									
RN56	0.949±0.002	0.901±0.004	1.833±0.062	1.352±0.040	Logistic	0.951±0.004	0.902±0.005	1.109±0.016	0.956±0.007
ShuffV2	0.949±0.003	0.899±0.004	2.322±0.089	1.686±0.056	RF	0.951±0.005	0.901±0.006	0.970±0.005	0.908±0.006
VGG16	0.951±0.003	0.899±0.004	1.531±0.068	1.130±0.025	HistGB	0.950±0.005	0.901±0.006	0.958±0.005	0.903±0.005
DLA	0.950±0.003	0.899±0.005	1.620±0.063	1.218±0.037	MLP	0.950±0.003	0.908±0.010	0.958±0.004	0.912±0.010
EffNet	0.949±0.004	0.900±0.005	1.929±0.067	1.366±0.033					
<i>Aggregation Methods</i>									
BL	0.951±0.003	0.899±0.004	1.531±0.068	1.130±0.025	BL	0.950±0.005	0.901±0.006	0.958±0.005	0.903±0.005
CR	0.948±0.005	0.908±0.006	1.506±0.054	1.194±0.029	CR	0.962±0.003	0.926±0.005	0.975±0.004	0.931±0.005
CM	0.988±0.002	0.971±0.003	2.089±0.096	1.570±0.050	CM	0.961±0.018	0.967±0.003	0.975±0.024	0.988±0.004
Wagg	0.948±0.003	0.898±0.005	1.292±0.028	1.032±0.014	Wagg	0.950±0.004	0.901±0.005	0.955±0.004	0.903±0.005
CSA	0.950±0.005	0.898±0.007	1.294±0.037	1.038±0.018	CSA	0.952±0.004	0.900±0.006	0.958±0.005	0.903±0.006
SACP	0.950±0.003	0.899±0.004	1.308±0.033	1.034±0.014	SACP	0.951±0.004	0.901±0.005	0.956±0.004	0.904±0.005
SACP++	0.949±0.003	0.898±0.004	1.281±0.027	1.028±0.011	SACP++	0.948±0.004	0.897±0.005	0.954±0.004	0.900±0.006

Table 2: Empirical coverage (Cov) and average prediction set length (Len) for CIFAR-10 and MNIST with $\alpha \in \{0.05, 0.1\}$:

sets. Inputs and outputs are standardized using their empirical mean and variance. In regression, we construct a uniform grid of 255 points between $\min_{1 \leq i \leq n} Y_i$ and $\max_{1 \leq i \leq n} Y_i$.

For **CSA**, we draw $M = 50$ random projections and perform 20 binary-search iterations to calibrate coverage. For **Wagg**, we conduct a grid search of 200 weight vectors over the $(K - 1)$ -simplex. For the **CM** and **CR** methods, which guarantee coverage of $1 - 2\alpha$, we set their nominal mis-coverage rate to $\frac{\alpha}{2}$ for a fair comparison. For **SACP++**, we perform a grid search over p to minimize Equation 27. We use $K = 7$ diverse base regressors, including linear models, tree-based methods, neural networks, and Bayesian regressors. We report the average prediction-set length and empirical coverage on the test set: for regression datasets at $\alpha = 0.05$, and for classification datasets at both $\alpha = 0.05$ and $\alpha = 0.10$. Further implementation details are provided in Appendix B, and additional experiments, including an analysis of sensitivity to the number of predictors K , are presented in Appendix C.

Results and Discussion. The results across all datasets are summarized in Figure 2 and Table 2. In terms of coverage, our proposed methods, **SACP** and **SACP++**, consistently achieve the target empirical level across all datasets. Among baselines, **Wagg** also maintains coverage close to the nominal value, whereas **CSA** tends to under-cover, and **CM** and **CR** frequently exceed the nominal level, resulting in overly conservative (and thus unnecessarily large) prediction sets.

Regarding prediction-set length, our methods demonstrate strong overall efficiency. As expected, **SACP++** consistently yields shorter sets than **SACP**, while even the default **SACP** remains competitive with all baselines. For classification tasks, **SACP++** reliably produces the smallest prediction sets among all compared methods. In particular, on **CIFAR-10**, it achieves the lowest variance in prediction-set length across test samples—significantly smaller than that of both the base learners and other aggregation approaches. For regression tasks, **SACP++** outperforms the best individual model (**BL**) on five out of nine datasets and attains the top

overall performance among aggregation methods on seven out of nine datasets. These results highlight the benefit of aggregation in this setting: by leveraging shared structure across base predictors, SACP-based methods improve efficiency while maintaining valid coverage.

6 Conclusion

We proposed SACP, a novel method for aggregating non-conformity scores from K predictors to construct a single conformal prediction set. Our approach transforms calibration and test raw scores into e-variables, which are combined through symmetric functions. The method applies to any choice of a symmetric aggregation function, while its enhanced variant, SACP++, adaptively selects the function yielding the smallest prediction sets. For strictly monotonic symmetric aggregators, we consider the case when aggregation reduces to a sum of scalar functions, which facilitates both theoretical analysis, via a worst-case bound, and efficient numerical implementation through data-driven aggregation. Empirical results demonstrate that SACP is highly competitive, outperforming state-of-the-art conformal aggregation methods and even the best individual base learner in terms of prediction-set length across all classification datasets and in five out of nine regression datasets. These findings highlight the effectiveness of conformal aggregation in leveraging shared structure among predictors to improve efficiency while maintaining valid coverage. Future work includes extending SACP++ by exploring symmetric neural architectures capable of learning the optimal aggregation function directly, and investigating how the dependencies among predictors’ uncertainties influence overall performance and efficiency.

References

Abdar, M.; Pourpanah, F.; Hussain, S.; Rezazadegan, D.; Liu, L.; Ghavamzadeh, M.; Fieguth, P.; Cao, X.; Khosravi, A.; Acharya, U. R.; Makarenkov, V.; and Nahavandi, S.

2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76: 243–297.
- Angelopoulos, A. N.; and Bates, S. 2022. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. arXiv:2107.07511.
- Balinsky, A. A.; and Balinsky, A. D. 2024. Enhancing conformal prediction using e-test statistics. *arXiv preprint arXiv:2403.19082*.
- Blanchet, J.; Lam, H.; Liu, Y.; and Wang, R. 2024. Convolution bounds on quantile aggregation. *Operations Research*.
- Braun, S.; Aolaritei, L.; Jordan, M. I.; and Bach, F. 2025. Minimum Volume Conformal Sets for Multivariate Regression. arXiv:2503.19068.
- Breiman, L. 1996. Bagging predictors. *Machine learning*, 24(2): 123–140.
- Dheur, V.; Fontana, M.; Estievenart, Y.; Desobry, N.; and Ben Taieb, S. 2025. A unified comparative study with generalized conformity scores for multi-output conformal regression. In *The 42nd International Conference on Machine Learning*.
- Dietterich, T. G. 2000. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, 1–15. Springer.
- Gasparin, M.; and Ramdas, A. 2024a. Conformal online model aggregation. arXiv:2403.15527.
- Gasparin, M.; and Ramdas, A. 2024b. Merging uncertainty sets via majority vote. arXiv:2401.09379.
- Gasparin, M.; and Ramdas, A. 2025. Improving the statistical efficiency of cross-conformal prediction. *arXiv preprint arXiv:2503.01495*.
- Gasparin, M.; Wang, R.; and Ramdas, A. 2025. Combining exchangeable p-values. *Proceedings of the National Academy of Sciences*, 122(11): e2410849122.
- Gauthier, E.; Bach, F.; and Jordan, M. I. 2025a. Backward Conformal Prediction. *arXiv preprint arXiv:2505.13732*.
- Gauthier, E.; Bach, F.; and Jordan, M. I. 2025b. E-values expand the scope of conformal prediction. *arXiv preprint arXiv:2503.13050*.
- Jin, Y.; and Candès, E. J. 2023. Selection by Prediction with Conformal p-values. arXiv:2210.01408.
- Klein, M.; Bethune, L.; Ndiaye, E.; and Cuturi, M. 2025. Multivariate Conformal Prediction using Optimal Transport. arXiv:2502.03609.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- LeCun, Y.; Cortes, C.; and Burges, C. 2010. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2.
- Lichtendahl Jr, K. C.; Grushka-Cockayne, Y.; and Winkler, R. L. 2013. Is it better to average probabilities or quantiles? *Management Science*, 59(7): 1594–1611.
- Luo, R.; and Zhou, Z. 2024. Trustworthy Classification through Rank-Based Conformal Prediction Sets. arXiv:2407.04407.
- Luo, R.; and Zhou, Z. 2025. Weighted Aggregation of Conformity Scores for Classification. arXiv:2407.10230.
- Ramdas, A.; and Wang, R. 2024. Hypothesis testing with e-values. *arXiv preprint arXiv:2410.23614*.
- Rivera, E. O.; Patel, Y.; and Tewari, A. 2025. Conformal Prediction for Ensembles: Improving Efficiency via Score-Based Aggregation. arXiv:2405.16246.
- Thomas, E. A.; and Ross, B. H. 1980. On appropriate procedures for combining probability distributions within the same family. *Journal of Mathematical Psychology*, 21(2): 136–152.
- Thurin, G.; Nadjahi, K.; and Boyer, C. 2025. Optimal Transport-based Conformal Prediction. arXiv:2501.18991.
- Vanschoren, J.; van Rijn, J. N.; Bischl, B.; and Torgo, L. 2014. OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2): 49–60.
- Vincent, S. B. 1912. *The Functions of the Vibrissae in the Behavior of the White Rat...*, volume 1. University of Chicago.
- Vovk, V. 2025. Conformal e-prediction. *arXiv preprint arXiv:2001.05989*.
- Vovk, V.; Gammerman, A.; and Shafer, G. 2005. *Algorithmic learning in a random world*. Springer.
- Vovk, V.; Wang, B.; and Wang, R. 2022. Admissible ways of merging p-values under arbitrary dependence. *The Annals of Statistics*, 50(1): 351–375.
- Vovk, V.; and Wang, R. 2020. Combining p-values via averaging. *Biometrika*, 107(4): 791–808.
- Vovk, V.; and Wang, R. 2021. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3): 1736–1754.
- Wang, R. 2024. The only admissible way of merging e-values. *arXiv preprint arXiv:2409.19888*.
- Wang, R.; and Ramdas, A. 2022. False discovery rate control with e-values. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(3): 822–852.
- Wang, T.; Wang, Y.; Zhou, J.; Peng, B.; Song, X.; Zhang, C.; Sun, X.; Niu, Q.; Liu, J.; Chen, S.; Chen, K.; Li, M.; Feng, P.; Bi, Z.; Liu, M.; Zhang, Y.; Fei, C.; Yin, C. H.; and Yan, L. K. 2025. From Aleatoric to Epistemic: Exploring Uncertainty Quantification Techniques in Artificial Intelligence. arXiv:2501.03282.
- Yang, Y.; and Kuchibhotla, A. K. 2024. Selection and Aggregation of Conformal Prediction Sets. arXiv:2104.13871.
- Zaheer, M.; Kottur, S.; Ravanbakhsh, S.; Póczos, B.; Salakhutdinov, R.; and Smola, A. 2018. Deep Sets. arXiv:1703.06114.