

Localized conformal model selection

Yuhao Wang^{1,2} and Tengyao Wang³

¹ Tsinghua University

² Shanghai Qizhi Institute

³ London School of Economics and Political Science

Abstract

We propose a localized conformal model selection framework that integrates local adaptivity with post-selection validity for distribution-free prediction. By performing model selection symmetrically across calibration points using upper and lower surrogate intervals, we construct a data-dependent safe index set that contains the oracle model and preserves exchangeability. The resulting ensemble procedure retains exact finite-sample marginal coverage while adapting to spatial heterogeneity and model complexity. Simulations demonstrate substantial reductions in interval length compared to the best fixed model, especially in heterogeneous and low-noise settings.

1 Introduction

Uncertainty quantification is a central problem in modern statistical learning. Among available approaches, conformal prediction has emerged as a particularly attractive framework (Vovk, Gammerman and Shafer, 2005; Shafer and Vovk, 2008; Lei et al., 2018; Angelopoulos and Bates, 2023). It provides finite-sample, distribution-free prediction intervals under minimal assumptions. These guarantees hold regardless of model complexity, making conformal prediction applicable to high-dimensional regression, nonparametric smoothing, and modern machine learning methods such as random forests and neural networks (Romano, Patterson and Candes, 2019; Barber et al., 2021b).

In this paper, we study conformal prediction for regression problems in which the goal is to construct a prediction interval for a future response at a new covariate value, based on previously observed data drawn from an unknown distribution. In many practical settings, rather than relying on a single predictive model, practitioners consider a collection of candidate regression estimators that differ in smoothness, complexity, or modelling assumptions. Selecting or combining these models in a data-driven way is often crucial for achieving good predictive performance.

However, standard conformal prediction methods are designed for a fixed predictive model and do not directly account for the additional uncertainty introduced by model selection. When the same data are used both to select a model and to construct prediction intervals, the exchangeability assumptions underpinning conformal validity are typically violated. As a result, naively applying conformal prediction after model selection often leads to systematic undercoverage, a phenomenon that has been widely observed in practice (Liang, Zhu and Barber, 2024; Hegazy et al., 2025).

Recent work has begun to address this challenge. Yang and Kuchibhotla (2025) and Liang, Zhu and Barber (2024) develop frameworks for conformal inference after model selection that preserves finite-sample coverage guarantees without sample splitting by correcting the selection bias inherent in naive procedures.

Related ideas also appear in the growing literature on selective conformal inference, which studies post-selection validity under various data-dependent selection rules (Jin and Ren, 2025). While these approaches successfully control selection-induced bias, they operate at a global or population level and do not allow the selected model to vary across the covariate space. This limitation can be restrictive in nonparametric settings where the underlying regression function exhibits spatially heterogeneous structure.

A complementary line of work focuses on local adaptivity. Localized Conformal Prediction, introduced by Guan (2023), modifies classical conformal prediction by weighting calibration residuals according to their proximity to the test point, yielding prediction intervals that adapt to local noise levels and model fit and can substantially reduce interval length in regions with simpler structure. Recent extensions explore related localization ideas, where local weights or randomization are used to achieve improved local coverage properties beyond marginal guarantees (Hore and Barber, 2025). More broadly, a growing literature investigates conditional or approximately conditional coverage guarantees, highlighting both the possibilities and fundamental limitations of achieving coverage that adapts to covariates (Barber et al., 2021a; Gibbs, Cherian and Candès, 2025). However, these localized methods assume a fixed predictive model and do not directly address the additional uncertainty arising from model selection or aggregation.

The goal of this paper is to bridge these two strands of work by developing a conformal prediction framework that is both locally adaptive and valid after model selection. A naive approach would be to select, at each test point, the model that yields the shortest localized conformal interval. However, such a strategy breaks the symmetry required for conformal validity, as the model selection event depends implicitly on the unobserved response at the test point. Our main contribution is a new localized conformal model selection procedure that combines LCP with a safe indexing technique to restore the required symmetry. The key insight is that model selection must be performed symmetrically across all data points, rather than only at the final test point, in order to preserve exchangeability. By explicitly accounting for the unobserved response through upper and lower surrogate conformal intervals, we construct a data-dependent safe index set that is guaranteed to contain the oracle model, defined as the model yielding the shortest expected conformal interval.

This construction leads to a principled ensemble conformal prediction method that adapts both to local structure in the covariates and to model uncertainty, while retaining exact finite-sample coverage guarantees. The resulting prediction intervals automatically favor more complex models in highly oscillatory regions and simpler models in smoother regions, as illustrated in our numerical experiments. To the best of our knowledge, this is the first conformal inference framework that simultaneously achieves local adaptivity and post-selection validity in a fully nonparametric setting.

2 A new localized conformal model selection

We consider a regression setting in which we observe a calibration dataset

$$\mathcal{D} := \{(X_i, Y_i)\}_{i=1}^n,$$

consisting of independent and identically distributed observations from an unknown joint distribution on $\mathcal{X} \times \mathbb{R}$. Our goal is to construct a prediction interval for an unobserved response Y_{n+1} associated with a new covariate value X_{n+1} . Throughout, we assume that the augmented sample $\{(X_i, Y_i)\}_{i=1}^{n+1}$ is exchangeable. We are given a finite collection of candidate regression estimators f_1, \dots, f_K , obtained from an independent training sample and are treated as fixed functions.

Since multiple regression estimators $\{f_k\}_{k=1}^K$ are available, an important question is how to combine them in order to construct shorter prediction intervals. We address this by employing *localized con-*

formal prediction (LCP) (Guan, 2023) and by introducing a safe-indexing refinement for model selection. Unlike the classical conformal prediction which constructs prediction sets via using the empirical quantile of regression residuals, the main idea of LCP is to produce a prediction interval for X_{n+1} based on the estimated regression residuals $Y_i - f(X_i)$ and a localizer $H(\cdot, \cdot)$. Here localizer is used to up-weight the residuals whose X_i are close to X_{n+1} ; in other words, LCP constructed prediction intervals based on a “weighted” empirical quantiles, where the weight of each sample is decided based on how close the sample covariance is to X_{n+1} . As a result, the resulting prediction interval adapts to local noise levels and local model fit, rather than reflecting only global behaviour. Some typical localizers include the exponential kernel $H(X_{n+1}, X_i) = \exp(-\|X_{n+1} - X_i\|_2/h)$ and the Gaussian kernel $H(X_{n+1}, X_i) = \exp(-\|X_{n+1} - X_i\|_2^2/h^2)$ for suitable choice of localizer bandwidths h .

Suppose we are given K conditional expectation estimates f_1, \dots, f_K (from a separate training sample) and a localizer H , then for each f_k , we may construct a conformalized prediction interval $C_k(\cdot)$ for f_k based weighted empirical quantile of regression residuals, in a similar way as described in Guan (2023). More specifically, given a calibration set $\{(Y_i, X_i)\}_{i=1}^n$ that is independent from the trained conditional expectation estimates f_1, \dots, f_K , a test covariate X and a model-specific miscoverage level γ , we may define $C_k(\{(Y_i, X_i)\}_{i=1}^n, X, \gamma)$ as

$$C_k(\{(Y_i, X_i)\}_{i=1}^n, X, \gamma) := [f_k(X) - q, f_k(X) + q],$$

$$\text{where } q := \min \left\{ v : \sum_{i=1}^n \frac{H(X_i, X)}{\sum_{i=1}^n H(X_i, X)} \mathbb{1}\{|Y_i - f_k(X_i)| \leq v\} \geq 1 - \gamma \right\}. \quad (1)$$

If the miscoverage level γ is fixed in advance and a single model were used, this construction recovers standard localized conformal prediction. The challenge addressed in this paper is how to select or aggregate models without violating finite-sample conformal validity. For $i \in [n]$, define the leave-one-out dataset $\mathcal{D}_{-i} := \mathcal{D} \setminus \{(Y_i, X_i)\}$. Under the exchangeability assumption, it is natural to consider the conformalized prediction interval $C_k(\mathcal{D}_{-i} \cup \{(Y_{n+1}, X_{n+1})\}, X_i, \gamma)$ obtained by interchanging the roles of indices i and $n + 1$. Note that this is an “oracle interval” in the sense that depends on the unobserved response Y_{n+1} . In practice, we instead construct upper and lower surrogate intervals that correspond respectively to the hypothetical cases of extreme nonconformity and perfect conformity of Y_{n+1} :

$$C_k^+(\mathcal{D}_{-i}, X_i, \gamma) := C_k(\mathcal{D}_{-i} \cup \{(\infty, X_{n+1})\}, X_i, \gamma)$$

$$C_k^-(\mathcal{D}_{-i}, X_i, \gamma) := C_k(\mathcal{D}_{-i} \cup \{(f_k(X_{n+1}), X_{n+1})\}, X_i, \gamma).$$

For each $i \in [n]$, we define the *safe index set*

$$\mathcal{K}_i(\gamma) := \left\{ k : |C_k^-(\mathcal{D}_{-i}, X_i, \gamma)| \leq \min_{k'} |C_{k'}^+(\mathcal{D}_{-i}, X_i, \gamma)| \right\}.$$

This set contains all models whose best-case interval length is no larger than the worst-case interval length of the shortest competing model. By construction,

$$C_k^-(\mathcal{D}_{-i}, X_i, \gamma) \subseteq C_k(\mathcal{D}_{-i} \cup \{(X_{n+1}, Y_{n+1})\}, X_i, \gamma) \subseteq C_k^+(\mathcal{D}_{-i}, X_i, \gamma),$$

for every k . Hence any model minimizing the oracle interval length must belong to $\mathcal{K}_i(\gamma)$.

Using the safe index sets $\mathcal{K}_i(\gamma)$, we determine lower and upper bounds on the admissible miscoverage level by evaluating empirical coverage across all calibration points. Specifically, over a discrete grid $\Gamma \subseteq$

$[0, 1]$, we define

$$\begin{aligned}\hat{\gamma}_L(\alpha) &:= \max \left\{ \gamma \in \Gamma : \frac{1}{n+1} \sum_{i=1}^n \mathbb{1} \{Y_i \in \cap_{k \in \mathcal{K}_i(\gamma)} C_k^-(\mathcal{D}_{-i}, X_i, \gamma)\} \geq 1 - \alpha \right\}, \\ \hat{\gamma}_U(\alpha) &:= \max \left\{ \gamma \in \Gamma : \frac{1}{n+1} \sum_{i=1}^n \mathbb{1} \{Y_i \in \cup_{k \in \mathcal{K}_i(\gamma)} C_k^+(\mathcal{D}_{-i}, X_i, \gamma)\} \geq 1 - \alpha - \frac{1}{n+1} \right\}.\end{aligned}\quad (2)$$

Finally, the prediction interval for Y_{n+1} is obtained by selecting, for each admissible miscoverage level, the shortest localized conformal interval and taking their union. Formally,

$$C_{\text{LCP-MS}}(\mathcal{D}, X_{n+1}, \alpha) := \bigcup_{\gamma \in [\hat{\gamma}_L(\alpha), \hat{\gamma}_U(\alpha)] \cap \Gamma} C_{\min}(\mathcal{D}, X_{n+1}, \gamma), \quad (3)$$

where $C_{\min}(\mathcal{D}, X_{n+1}, \gamma)$ denotes the shortest interval among $\{C_k(\mathcal{D}, X_{n+1}, \gamma)\}_{k=1}^K$, with ties broken by index order.

The following theorem establishes the finite-sample validity of the resulting procedure.

Theorem 1. *Consider a dataset $\{(Y_i, X_i)\}_{i=1}^{n+1}$ that are i.i.d. realizations from a distribution \mathbb{P} , where Y_{n+1} is unobserved. Given a sequence of fixed functions f_1, \dots, f_K , we have that the $C_{\text{LCP-MS}}(\cdot)$ constructed in (3) satisfies*

$$\mathbb{P}(Y_{n+1} \in C_{\text{LCP-MS}}(\mathcal{D}, X_{n+1}, \alpha)) \geq 1 - \alpha$$

for any $\alpha \in [0, 1]$.

Proof. Let $\tilde{\mathcal{D}} := \mathcal{D} \cup \{(X_{n+1}, Y_{n+1})\}$ and $\tilde{\mathcal{D}}_{-i} := \tilde{\mathcal{D}} \setminus \{(X_i, Y_i)\}$. Define

$$\hat{\gamma} := \max \left\{ \gamma \in \Gamma : \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{1} \{Y_i \in C_{\min}(\tilde{\mathcal{D}}_{-i}, X_i, \gamma)\} \geq 1 - \alpha \right\}. \quad (4)$$

Set

$$Z_i := \mathbb{1} \{Y_i \in C_{\min}(\tilde{\mathcal{D}}_{-i}, X_i, \hat{\gamma})\}.$$

By exchangeability of $\{(X_i, Y_i)\}_{i=1}^{n+1}$ and the permutation invariance of the construction, the variables Z_1, \dots, Z_{n+1} are exchangeable. By definition of $\hat{\gamma}$, and using exchangeability, we have

$$\mathbb{P}(Y_{n+1} \in C_{\min}(\mathcal{D}, X_{n+1}, \hat{\gamma})) = \mathbb{P}(Z_{n+1} = 1) = \mathbb{E} \left[\frac{1}{n+1} \sum_{i=1}^{n+1} Z_i \right] \geq 1 - \alpha.$$

Thus, it remains to show that $\hat{\gamma}_L(\alpha) \leq \hat{\gamma} \leq \hat{\gamma}_U(\alpha)$ for $\hat{\gamma}_L(\alpha)$ and $\hat{\gamma}_U(\alpha)$ defined in (2).

To this end, from the surrogate inclusions, we have for all $i \in [n]$ and $\gamma \in \Gamma$ that

$$C_k^-(\mathcal{D}_{-i}, X_i, \gamma) \subseteq C_k(\tilde{\mathcal{D}}_{-i}, X_i, \gamma) \subseteq C_k^+(\mathcal{D}_{-i}, X_i, \gamma).$$

Hence the model f_k whose corresponding interval $C_k(\tilde{\mathcal{D}}_{-i}, X_i, \gamma)$ is shortest among all $C_{k'}(\tilde{\mathcal{D}}_{-i}, X_i, \gamma)$ for $k \in \{1, \dots, K\}$ must belong to the set $\mathcal{K}_i(\gamma)$. Consequently, for all $i \in [n]$, we have

$$\cap_{k \in \mathcal{K}_i(\gamma)} C_k^-(\mathcal{D}_{-i}, X_i, \gamma) \subseteq C_{\min}(\tilde{\mathcal{D}}_{-i}, X_i, \gamma) \subseteq \cup_{k \in \mathcal{K}_i(\gamma)} C_k^+(\mathcal{D}_{-i}, X_i, \gamma).$$

Then it follows from above, (4), as well as the definitions of $\hat{\gamma}_L(\alpha)$ and $\hat{\gamma}_U(\alpha)$ in (2) that $\hat{\gamma}_L(\alpha) \leq \hat{\gamma} \leq \hat{\gamma}_U(\alpha)$, which proves the desired result. \square

3 Numerical simulations

We illustrate the proposed localized conformal model selection (LCP-MS) method through two simulation studies: a parametric setting with locally oscillatory signals and a nonparametric regression setting with varying curvature. In all experiments, we report the average conformal interval length at the target coverage level $1 - \alpha = 0.9$ based on 100 replications. We compare the proposed LCP-MS ensemble with the best-performing single model chosen in hindsight (denoted “best_single”).

We first consider a piecewise sinusoidal model defined as

$$Y = \begin{cases} \sin(5X) + \varepsilon, & X \in [-5, 0), \\ 2 \sin(3X) + \varepsilon, & X \in [0, 5], \end{cases} \quad \varepsilon \sim N(0, \sigma^2),$$

where the regression function changes its frequency and amplitude at $X = 0$.

We fit a family of local parametric models

$$M_{\lambda, h} : Y = A \sin(\lambda X + \phi),$$

where the amplitude A and phase shift ϕ are estimated via least squares using only data within a local window $[x - h, x + h]$. Each model is indexed by a frequency parameter $\lambda \in \{1, 2, 3, 4, 5\}$ and a local bandwidth $h \in \{0.5, 1\}$, resulting in 10 candidate models in total.

We vary the noise level $\sigma \in \{0.1, 0.3\}$ and the training/calibration sample size $n \in \{200, 500, 1000\}$. For each configuration, Table 1 reports the average length of the conformal prediction intervals obtained by the proposed LCP-MS ensemble and by the best individual model.

Table 1: Average conformal interval lengths in the parametric setting.

n	σ	localizer bw	ensemble_len	best_single
200	0.1	0.1	0.3985	0.4865
500	0.1	0.1	0.3660	0.4445
1000	0.1	0.1	0.3470	0.4305
200	0.3	0.1	1.0920	1.1020
500	0.3	0.1	1.0950	1.0460
1000	0.3	0.1	1.0455	1.0330
200	0.1	0.3	0.3840	0.4935
500	0.1	0.3	0.3550	0.4460
1000	0.1	0.3	0.3450	0.4295
200	0.3	0.3	1.1620	1.1130
500	0.3	0.3	1.0475	1.0480
1000	0.3	0.3	1.0155	1.0385

Across all sample sizes, the proposed LCP-MS ensemble consistently produces shorter intervals than the best single model under low noise ($\sigma = 0.1$), demonstrating its benefit from adaptive local model selection. Under higher noise ($\sigma = 0.3$), both methods perform similarly, indicating that the advantage of localization diminishes when the signal-to-noise ratio is low.

We next consider a smooth nonparametric regression model

$$Y = \sin(X^3) + \varepsilon, \quad X \in [0, 3], \quad \varepsilon \sim N(0, \sigma^2),$$

where the curvature of the mean function varies with X . We fit five Nadaraya–Watson kernel regression estimators with Gaussian kernels and bandwidths $h \in \{0.1, 0.2, 0.4, 0.8, 1.6\}$. The same LCP-MS procedure is applied to combine these models.

We vary $\sigma \in \{0.1, 0.3\}$ and $n \in \{200, 500, 1000, 2000\}$. Table 2 summarizes the average interval lengths. In this nonlinear setting, the LCP-MS intervals are substantially shorter, roughly 25–35% reduction compared with the best single model, while maintaining nominal coverage. The improvement is more pronounced for small σ and large n , highlighting the ability of the proposed method to adapt to local structure and select the most appropriate kernel bandwidths across regions of differing smoothness.

Table 2: Average conformal interval lengths in the nonparametric setting.

n	σ	localizer bw	ensemble_len	best_single
200	0.1	0.3	0.9500	1.3665
500	0.1	0.3	0.9245	1.3265
1000	0.1	0.3	0.9045	1.3175
2000	0.1	0.3	0.9000	1.3225
200	0.3	0.3	1.4580	1.5790
500	0.3	0.3	1.4145	1.5495
1000	0.3	0.3	1.4170	1.5530
2000	0.3	0.3	1.3960	1.5390

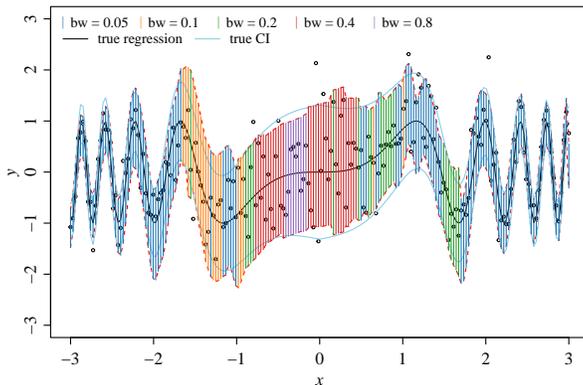


Figure 1: Localized ensemble conformal prediction intervals evaluated at each test point. The black curve denotes the true mean regression function, while the light blue curves indicate the true 90% confidence bands. Vertical bars represent ensemble conformal prediction intervals, with colours corresponding to different models selected within the ensemble. In highly oscillatory regions, models with smaller bandwidths are predominantly selected (blue), whereas in smoother regions the ensemble favours models with larger bandwidths (red and purple).

To illustrate the adaptive behaviour underlying the improvements shown in the simulation results, we focus on a representative configuration with $n = 500$, $\sigma = 0.1$ and $h = 0.3$ and examine the resulting localized ensemble conformal intervals in Figure 1. The figure shows that the LCP-MS ensemble adaptively selects different kernel bandwidths across the covariate domain in accordance with the local geometry of the true regression function. In regions where the curvature is high, smaller bandwidth models are pre-

dominantly chosen, producing tighter intervals that closely track rapid changes in the signal. In contrast, in smoother regions the ensemble favours larger bandwidths, yielding more stable and efficient inference. This spatially varying model selection provides a clear explanation for the superior performance of the ensemble observed in Table 2, as no single fixed-bandwidth estimator can adapt uniformly well across regions of differing smoothness.

4 Discussion

The main technical challenge addressed in the paper is to reconcile local, data dependent model selection with the symmetry requirements of conformal inference. Selecting the shortest localized interval at a test point alone breaks exchangeability. The safe indexing construction resolves this by enforcing selection rules simultaneously across calibration points through surrogate intervals.

We remark that while the procedure improves interval adaptivity across the covariate space, the overall guarantee remains marginal rather than conditional. Efficiency gains are most pronounced when the candidate family contains a model that approximates the local regression function well. In such settings, and when the best and second best models are well separated, the safe index set $\mathcal{K}_i(\gamma)$ is expected to become asymptotically a singleton, so the method effectively behaves like a locally optimal selector.

The procedure is computationally heavier than standard conformal prediction due to leave one out constructions and evaluation over multiple models and miscoverage levels, though these computations are naturally parallelizable. While developed for regression, the symmetry based safe indexing idea may extend to other adaptive conformal settings where data dependent decisions threaten exchangeability.

References

- Angelopoulos, A. N. and Bates, S. (2023) Conformal prediction: A gentle introduction. *Found. Trends Mach. Learn.*, **16**, 494–591.
- Barber, R. F., Candes, E. J., Ramdas, A. and Tibshirani, R. J. (2021a) The limits of distribution-free conditional predictive inference. *Inf. Inference*, **10**, 455–482.
- Barber, R. F., Candes, E. J., Ramdas, A. and Tibshirani, R. J. (2021b) Predictive inference with the jackknife+. *Ann. Statist.*, **49**, 486–507.
- Gibbs, I., Cherian, J. J. and Candès, E. J. (2025) Conformal prediction with conditional guarantees. *J. Roy. Statist. Soc., Ser. B*, **87**, 1100–1126.
- Guan, L. (2023) Localized conformal prediction: A generalized inference framework for conformal prediction. *Biometrika*, **110**, 33–50.
- Hegazy, M., Aolaritei, L., Jordan, M. I. and Dieuleveut, A. (2025) Valid Selection among Conformal Sets. *Adv. Neural Inf. Process. Syst.*, **38**.
- Hore, R. and Barber, R. F. (2025) Conformal prediction with local weights: randomization enables local guarantees. *J. Roy. Statist. Soc., Ser. B*, **87**, 549–578.
- Jin, Y. and Ren, Z. (2025) Confidence on the focal: Conformal prediction with selection-conditional coverage. *J. Roy. Statist. Soc., Ser. B*, **87**, 1239–1259.

- Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R. J. and Wasserman, L. (2018) Distribution-free predictive inference for regression. *J. Amer. Statist. Assoc.*, **113**, 1094–1111.
- Liang, R., Zhu, W. and Barber, R. F. (2024) Conformal prediction after data-dependent model selection. *arXiv preprint arXiv:2408.07066*.
- Romano, Y., Patterson, E. and Candes, E. (2019) Conformalized quantile regression. *Adv. Neur. Inform. Proc. Syst.*, **32**.
- Shafer, G. and Vovk, V. (2008) A tutorial on conformal prediction. *J. Mach. Learn. Res.*, **9**.
- Vovk, V., Gammerman, A. and Shafer, G. (2005) *Algorithmic Learning in a Random World*. Springer.
- Yang, Y. and Kuchibhotla, A. K. (2025) Selection and aggregation of conformal prediction sets. *J. Amer. Statist. Assoc.*, **120**, 435–447.