

Resolving Predictive Multiplicity for the Rashomon Set

Parian Haghghat¹, Hadis Anahideh¹, Cynthia Rudin²

¹University of Illinois Chicago, Chicago, IL, USA

²Duke University, Durham, NC, USA

Abstract

The existence of multiple, equally accurate models for a given predictive task leads to predictive multiplicity, where a “Rashomon set” of models achieve similar accuracy but diverge in their individual predictions. This inconsistency undermines trust in high-stakes applications where we want consistent predictions. We propose three approaches to reduce inconsistency among predictions for the members of the Rashomon set. The first approach is **outlier correction**. An outlier has a label that none of the good models are capable of predicting correctly. Outliers can cause the Rashomon set to have high variance predictions in a local area, so fixing them can lower variance. Our second approach is **local patching**. In a local region around a test point, models may disagree with each other because some of them are biased. We can detect and fix such biases using a validation set, which also reduces multiplicity. Our third approach is **pairwise reconciliation**, where we find pairs of models that disagree on a region around the test point. We modify predictions that disagree, making them less biased. These three approaches can be used together or separately, and they each have distinct advantages. The reconciled predictions can then be distilled into a single interpretable model for real-world deployment. In experiments across multiple datasets, our methods reduce disagreement metrics while maintaining competitive accuracy.

Code — <https://github.com/parianh/ReconciledRashomon>

Introduction

The *Rashomon Effect* is when a dataset admits a diverse set of models that perform about equally well (McCullagh and Nelder 1989; Breiman 2001; Rudin et al. 2024). This gives rise to *predictive multiplicity*, where models with similar global accuracy have divergent individual-level predictions. This has been observed in credit scoring, criminal sentencing, and medical diagnosis, where multiple equally accurate models disagree on individual predictions (Marx et al. 2020; Black, Raghavan, and Solon 2022; Kobylińska et al. 2024).

Model multiplicity has both practical advantages and disadvantages. On the positive side, the diversity of equally accurate models reflects the inherent flexibility of many learning problems and opens the door to aligning system behavior with crucial desiderata such as fairness, robustness, and

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Illustration of approaches for regression.

interpretability (Rudin 2019; Black, Raghavan, and Solon 2022; Rudin et al. 2024). On the negative side, inconsistency in the predictions of equally good models potentially undermines trust in algorithmic decisions. Relying on only one good model risks reinforcing systemic blind spots and can leave decisions vulnerable to be challenged (Kleinberg et al. 2021; D’Amour et al. 2022).

In this work, we show how the Rashomon set, i.e., the set of good models for a given dataset, can be reconciled to agree on predictions in a way that does not harm accuracy. With each edit we make to a model in the Rashomon set, we lessen bias and improve local accuracy, which is why the model predictions iteratively converge toward agreement.

We propose three approaches for reconciling predictions from the Rashomon set: **outlier correction**, **local patching**, and **pairwise reconciliation**, illustrated in Figure 1; our algorithms are for classification, but it is easier to illustrate them for regression. These approaches can be used separately or together, and each has different benefits.

Outlier Correction (OC), as well as omitting outliers, is common as a data preprocessing step (Breunig et al. 2000; Liu, Ting, and Zhou 2012; Zhao, Nasrullah, and Li 2019) to reduce the bias of models. Here, we use it for a different purpose: to reduce the variance of predictions among members of the Rashomon set. Models trained with outliers are generally worse and run the risk of overfitting. We identify outliers as points whose labels disagree with the predictions of models in the Rashomon set.

We consider both full OC (where we correct outliers in training and validation sets, then retrain) and a lighter OC variant that only corrects validation set outliers (without retraining). As we show, outlier correction can improve predictive performance and stability, but it does not consistently reduce ensemble-wide disagreement on its own. We find that OC is most useful as a preprocessing step for subse-

quent reconciliation, in particular, correcting validation outlier labels experimentally provides cleaner data for our local patching method, discussed next.

In **local patching (LP)**, if a model exhibits systematic over- or under-prediction around a test point, we reduce this local bias by editing the model’s prediction at that point. Specifically, for each test point, we use a validation-set neighborhood to estimate directional bias for each model. We then patch predictions by adjusting them toward the neighborhood’s validation labels. This method empirically improves local calibration error and reduces local bias for models in the Rashomon set.

Our third approach is **pairwise reconciliation (PR)**. This method is inspired by the theoretical reconciliation framework of Roth et al. (2023), although our implementation addresses practical aspects that the theory does not. The general goal of reconciliation is to identify regions of systematic disagreement among accurate models and edit predictions to reduce inconsistency. Our approach differs from that of Roth et al. (2023) in important ways. Our reconciliation decisions take advantage of held-out validation data, which we use to assess empirical suboptimality and verify that any proposed edits improve predictive performance. In addition, while Roth et al. (2023) perform reconciliation only by correcting one model to be similar to another in regions where they disagree, we make corrections by trading off predictive accuracy against resolving disagreement between the models. Finally, rather than reconciling model pairs in an arbitrary order, we prioritize the pairs of models exhibiting the largest disagreement, which improves computational efficiency and accelerates convergence in practice.

In our pairwise reconciliation approach, (1) We iteratively identify the pair of models from the Rashomon set that disagree the most on predictions for the validation set; (2) For each pair, we determine in which direction they most disagree (either the first model is larger more often or the second one is). The set of points on which the models disagree by at least ϵ in this direction is called the *validation disagreement set*. (3) We determine which of the two models is empirically suboptimal on the validation disagreement set – i.e., has worse predictive performance on its validation labels – and we say that model is *falsified*. (4) We then apply a correction to the falsified model’s predictions on the validation disagreement set, where the correction magnitude is a linear combination of model error and disagreement between the models. (5) We transfer the correction to test points where the same model pair disagrees in the same direction as in the *validation disagreement set* used to compute that correction.

In this way, reconciliation edits predictions in regions of the test space where the disagreement structure is likely to persist. Empirically, pairwise reconciliation consistently and substantially reduces ensemble-wide disagreement across datasets, often driving variance, ambiguity, discrepancy, and disagreement-based metrics close to zero. However, reconciliation alone does not directly address cases where models share a common bias and therefore do not disagree—such as underfitting or systematic local miscalibration. This is where local patching and outlier correction help. Local patching improves neighborhood agreement by correcting directional

local bias, while outlier correction provides cleaner data for both local patching and reconciliation. We find that combining all three approaches yields robust reductions in disagreement while maintaining competitive predictive performance.

The reconciled predictions represent a consensus among good models; they are less locally biased than individual models and are more difficult to challenge than any single model. They have numerous possible options for downstream deployment; for instance, to support interpretability and accountability, we can distill the reconciled predictions into an interpretable meta-model that closely approximates the ensemble’s final behavior.

Related Work

Predictive Multiplicity and the Rashomon Set. Throughout the history of statistics, scientists have noted the Rashomon Effect and the phenomenon of predictive multiplicity — when many models achieve near-identical accuracy, but produce conflicting predictions for the same points (McCullagh and Nelder 1989; Breiman 2001; Rudin 2019; D’Amour et al. 2022; Marx et al. 2020; Black, Raghavan, and Solon 2022; Rudin et al. 2024). The set of such well-performing but diverse models is called the *Rashomon set* (Fisher, Rudin, and Dominici 2019; Semenova, Rudin, and Parr 2022), and recent work has focused on quantifying and characterizing the extent of disagreement among models within the Rashomon set. Marx et al. (2020) proposed a set of metrics to measure disagreement in classification tasks, which Watson-Daniels, Parkes, and Ustun (2023) extended to probabilistic predictions, highlighting that predictive variance is especially pronounced near decision boundaries and disproportionately affects individuals from minority groups. Coker, Rudin, and King (2021) defined tethered hacking intervals to measure the range of a statistic across the Rashomon set. If the statistic is a model prediction, the tethered hacking intervals will be prediction intervals that are robust to arbitrary choices made by scientists who might reasonably choose any model from the Rashomon set.

In parallel, advances in model enumeration techniques have enabled researchers to systematically explore the Rashomon set for specific model families. For example, complete enumerations have been developed for sparse decision trees (Xin et al. 2022), illustrating the sheer scale and structural diversity of high-performing models that can emerge from the same training data (Rudin et al. 2024).

Predictive multiplicity gives rise to what Black, Raghavan, and Solon (2022) term a “crisis of justifiability.” In high-stakes settings such as lending or criminal justice, if one model denies a loan while another—equally accurate—would approve it, the affected individual may have grounds to question the fairness and objectivity of the decision. Such inconsistencies at the individual level not only erode trust but also complicate contestability and recourse (Long et al. 2024).

Reconciliation Strategies. As discussed, Roth et al. (2023) propose a pairwise reconciliation framework that identifies regions of substantial disagreement between models and adjusts predictions to be closer on those disagreement regions. Our approach grounds decisions in held-out

validation data and makes corrections by trading off predictive accuracy against reduced disagreement between models. Our approach is computationally efficient in practice, since it prioritizes the pairs with the most extreme disagreement rather than reconciling pairs in an arbitrary order.

Du, Ngo, and Wu (2024) introduce a multi-calibration framework that reconciles predictions across groups by post-processing a base predictor (e.g., a score from an ensemble model) to be simultaneously calibrated on overlapping subpopulations. While this improves fairness metrics, it generally introduces additional bias that conflicts with our goals.

Approaches

We introduce our three approaches below. We are in a standard supervised binary classification setting, with i.i.d. points $\{x_i, y_i\}_{i=1}^n$, with $x \in \mathcal{X}$ and $y_i \in \{0, 1\}$. We distinguish training points, validation points, and test points with superscripts or subscripts. The training set is $\mathcal{D}^{\text{train}}$, the validation set is \mathcal{D}^{val} , and the test set is $\mathcal{D}^{\text{test}}$.

We assume we have access to the Rashomon set, or a collection of good models. Complete Rashomon sets can be computed for sparse decision trees (Xin et al. 2022) or linear models (which are ellipsoids around the least squares solution). A collection of good models can be constructed by running several different machine learning algorithms on the data; this is what we use in our experiments. The collection of models is denoted $f_m : \mathcal{X} \rightarrow [0, 1]$ for $m = 1..M$. The models predict the probability that $y = 1 \mid x$. We use $\text{clip}_{[0,1]}(p) = \min\{1, \max\{0, p\}\}$ to ensure corrected probabilities remain in $[0, 1]$.

We measure probabilistic prediction error using the **Brier loss**. For any finite dataset $\{(x_i, y_i)\}_{i=1}^n$ (e.g., the validation set \mathcal{D}^{val}) with $y_i \in \{0, 1\}$ and any index set $\mathcal{I} \subseteq \{1, \dots, n\}$, define the empirical Brier loss of a predictor f on \mathcal{I} as

$$\text{BrierLoss}_{\mathcal{I}}(f) := \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} (f(x_i) - y_i)^2. \quad (1)$$

When $\mathcal{I} = \{1, \dots, n_{\text{val}}\}$, this corresponds to the standard validation Brier score. Our setup can be generalized to multiclass classification or regression, as well as to infinite sets of functions.

Outlier Correction

Algorithm 1 shows how we correct outliers for classification, which improves label quality prior to reconciliation. Here, outliers are either true positives, where the average Rashomon set prediction is highly negative, or true negatives, where the average Rashomon set prediction is highly positive. For training outliers, we flip the label; this way, we can retrain the Rashomon set on the corrected labels afterwards if desired. For validation outliers, we replace the label with the mean predicted probability of being positive among the Rashomon set.

One can create an analogous algorithm for regression, where an outlier would be a point that no or few good models can correctly predict. For instance, if the label is greater than all of the Rashomon set’s predictions or lower than all of them, we would call it an outlier. We could also use a

Algorithm 1 Outlier Correction (for classification)

Require: Training predictions $f_m(x_i^{\text{train}})$ for each training point i and each model from the Rashomon set m , validation predictions $f_m(x_i^{\text{val}})$ for each validation point i and each model from the Rashomon set m (note that i indexes both training and validation), labels $y_i^{\text{train}}, y_i^{\text{val}}$ for each i , ensemble size M , lower/upper thresholds for the estimated probability $(\tau_{\text{low}}, \tau_{\text{high}})$.

- 1: Compute ensemble means for both training and validation points: $\hat{f}(x_i^{\text{train}}) = \frac{1}{M} \sum_m f_m(x_i^{\text{train}})$ and $\hat{f}(x_i^{\text{val}}) = \frac{1}{M} \sum_m f_m(x_i^{\text{val}})$.
- 2: Identify outliers using the thresholds $(\tau_{\text{low}}, \tau_{\text{high}})$: $\mathcal{O} = \{i : [(y_i^{\text{train or val}} = 1) \ \& \ (\hat{f}(x_i^{\text{train or val}}) < \tau_{\text{low}})] \text{ or } [(y_i^{\text{train or val}} = 0) \ \& \ (\hat{f}(x_i^{\text{train or val}}) > \tau_{\text{high}})]\}$ including points from both training and validation sets.
- 3: For each outlier i in the training set, flip its label: $y_i^{\text{train}} \leftarrow 1 - y_i^{\text{train}}$. For each outlier i in the validation set, set $y_i^{\text{val}} \leftarrow \hat{f}(x_i^{\text{val}})$.
- 4: If desired, recompute Rashomon set models using updated training labels y^{train} .
- 5: **return** Updated labels, cleaned training data, and new Rashomon set.

percentile (e.g., an outlier has a label above the 95th percentile or below the 5th percentile of predictions from the Rashomon set).

Local Patching

Algorithm 2 describes local patching. Here, we find local regions where a model disagrees with the labels in its neighborhood, and we issue a patch to fix it. This directly improves the local accuracy of each model by identifying and correcting local biases, or “blind spots.” For each test instance x^{test} , we first construct a context-specific neighborhood $N(x^{\text{test}}) \subset \mathcal{D}^{\text{val}}$ based on the radius of its k th nearest neighbor in the validation set. Within this neighborhood, each model f_m is evaluated for systematic local bias. The residuals for model f_m on validation neighbor j are $R_m = \{y_j^{\text{val}} - f_m(x_j^{\text{val}})\}$; they are computed using validation points in the neighborhood $N(x^{\text{test}})$. These residuals estimate the conditional error distribution $\varepsilon_m(x^{\text{test}}) = \mathbb{E}_{Y \mid x^{\text{test}}}[Y - f_m(x^{\text{test}})]$ from the data. Under the assumption of local smoothness, the expectation of residuals in $N(x^{\text{test}})$ serves as a consistent estimator for the bias in f_m ’s prediction at x^{test} .

To avoid overfitting to symmetric noise or canceling variance, we apply a one-sided bias test: we only compute correction terms when residuals exhibit directional skew above a threshold τ_{bias} . Formally, the patch Δ_m estimates the conditional bias $\varepsilon_m(x^{\text{test}})$ when it is sufficiently large. The patched prediction $f'_m(x^{\text{test}}) = f_m(x^{\text{test}}) + \Delta_m$ thus serves as a locally calibrated output.

$$\Delta_m = \begin{cases} \frac{1}{|R_m^+|} \sum_{r \in R_m^+} r & \text{if } \frac{|R_m^+|}{|R_m^-|} > \tau_{\text{bias}} \\ \frac{1}{|R_m^-|} \sum_{r \in R_m^-} r & \text{if } \frac{|R_m^-|}{|R_m^+|} > \tau_{\text{bias}} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where $R_m^+ = \{r \in R_m \mid r > 0\}$ and $R_m^- = \{r \in R_m \mid r < 0\}$ are the sets of positive and negative residu-

Algorithm 2 Local Patching

Require: Validation set \mathcal{D}^{val} , test set $\mathcal{D}^{\text{test}}$, model predictions for the validation set $\{f_m(x_i^{\text{val}})\}_{m=1..M, i=1..n_{\text{val}}}$, query point x^{test} and model predictions $\{f_m(x^{\text{test}})\}_{m=1}^M$, neighbor parameters (k, k_{max}) , bias threshold τ_{bias} .

- 1: Query the test point's k_{max} nearest neighbors within the validation set, $\text{KNN}(x^{\text{test}})$.
- 2: Filter the $\text{KNN}(x^{\text{test}})$ to form neighborhood $N(x^{\text{test}})$.
- 3: **for** each model $m = 1, \dots, M$ **do**
- 4: Compute residuals in $N(x^{\text{test}})$:

$$R_m(x^{\text{test}}) = \{y_j^{\text{val}} - f_m(x_j^{\text{val}}) : j \in N(x^{\text{test}})\}.$$

- 5: Compute the fractions of positive and negative residuals, and compute patch Δ_m using Equation 2. ▷ keep the patch only if it does not worsen neighborhood Brier loss (otherwise set $\Delta_m \leftarrow 0$)
 - 6: **if** $\Delta_m \neq 0$ **then**
 - 7: $b_{\text{before}} \leftarrow \frac{1}{|N(x^{\text{test}})|} \sum_{j \in N(x^{\text{test}})} (f_m(x_j^{\text{val}}) - y_j^{\text{val}})^2$
 - 8: $b_{\text{after}} \leftarrow \frac{1}{|N(x^{\text{test}})|} \sum_{j \in N(x^{\text{test}})} (\text{clip}_{[0,1]}(f_m(x_j^{\text{val}}) + \Delta_m) - y_j^{\text{val}})^2$
 - 9: **if** $b_{\text{after}} > b_{\text{before}}$ **then**
 - 10: $\Delta_m \leftarrow 0$ ▷ reject patch
 - 11: **end if**
 - 12: **end if**
 - 13: $f'_m(x^{\text{test}}) \leftarrow \text{clip}_{[0,1]}(f_m(x^{\text{test}}) + \Delta_m)$
 - 14: **end for**
 - 15: **return** $f'_m(x^{\text{test}})$ for all m .
-

als from $N(x^{\text{test}})$, respectively. $\tau_{\text{bias}} \in (0.5, 1]$ is a hyperparameter that controls sensitivity, and $|\cdot|$ is the count of elements in a set. To avoid harmful corrections, after computing Δ_m from the validation neighborhood $N(x^{\text{test}})$, we evaluate $\text{BrierLoss}_{N(x^{\text{test}})}(f_m)$ on the validation neighbors in $N(x^{\text{test}})$ before and after applying the (clipped) shift. We keep the patch at x^{test} only if this neighborhood Brier loss does not increase; otherwise, we set $\Delta_m = 0$. The patched prediction is then $f'_m(x^{\text{test}}) = \text{clip}_{[0,1]}(f_m(x^{\text{test}}) + \Delta_m)$, improving local calibration without worsening neighborhood Brier loss.

Pairwise Reconciliation

Roth et al. (2023) points out that when two models substantially disagree on a region of the input space, at least one must be statistically suboptimal on that region and can be corrected. Repeating this process iteratively yields a low-bias consensus. In our approach, shown in Algorithm 3, we iteratively identify pairs of models in the Rashomon set that exhibit the largest predictive disagreement on the validation set. We quantify pairwise disagreement on validation using ℓ_1 distance:

$$D[m, j] = \frac{1}{n_{\text{val}}} \sum_{i=1}^{n_{\text{val}}} |f_m(x_i^{\text{val}}) - f_j(x_i^{\text{val}})|,$$

and at each iteration, we select the B pairs (m, j) with the largest $D[m, j]$ values.

We define the ensemble consensus on validation as the

(fixed) initial mean prediction

$$\hat{f}^{\text{init}}(x_i^{\text{val}}) = \frac{1}{M} \sum_{m=1}^M f_m(x_i^{\text{val}}),$$

computed once before iterations.

For each selected pair, we form the signed disagreement sets $\mathcal{S}^>$ and $\mathcal{S}^<$ (defined below) consisting of validation points where the two models differ by more than ε , and we select one of them (formalized below), denoted \mathcal{S} .

We then identify which of the two models performs worse on that set using validation labels; we say that this model is *falsified* on \mathcal{S} . We then apply a correction to the falsified model's predictions on this set, shifting them toward the observed labels as well as the ensemble's consensus \hat{f}^{init} . We repeat this procedure until (i) the maximum pairwise disagreement falls below η , (ii) no further verified corrections are accepted (i.e., none satisfy the acceptance test), or (iii) the maximum number of iterations is reached.

By repeatedly reconciling the most disagreeing model pairs, we reduce disagreement across the ensemble while preserving accuracy.

There are several major differences between pairwise reconciliation and local patching. They have different goals: PR's goal is to encourage predictions of models to agree, whereas LP's goal is to reduce bias. Thus, LP considers individual models, whereas pairwise reconciliation considers pairs of models. LP applies a patch uniformly to a local region, whereas pairwise reconciliation applies corrections only at (or near) points where this same pair of models disagrees.

For each selected model pair (m, j) , we construct two *signed disagreement sets* from the validation set:

$$\mathcal{S}^> = \{i \in \{1, \dots, n_{\text{val}}\} : f_m(x_i^{\text{val}}) - f_j(x_i^{\text{val}}) > \varepsilon\}, \quad (3)$$

$$\mathcal{S}^< = \{i \in \{1, \dots, n_{\text{val}}\} : f_j(x_i^{\text{val}}) - f_m(x_i^{\text{val}}) > \varepsilon\}. \quad (4)$$

These sets capture points where one model exceeds the other by more than the disagreement threshold ε . We then choose the larger set and record its sign:

$$\mathcal{S} \leftarrow \arg \max_{\mathcal{A} \in \{\mathcal{S}^>, \mathcal{S}^<\}} |\mathcal{A}|,$$

$$\sigma \in \{>, <\} \text{ indicates whether } \mathcal{S} = \mathcal{S}^> \text{ or } \mathcal{S}^<.$$

(If there is a tie, we take $\mathcal{S} = \mathcal{S}^>$, matching Algorithm 3.)

If $|\mathcal{S}| < \alpha$, we skip this pair. Otherwise, let k denote the falsified model, i.e., the one with higher $\text{BrierLoss}_{\mathcal{S}}$ over \mathcal{S} (see Eq. (1)).

We then compute the correction z^* by minimizing the empirical convex objective over the selected disagreement index set \mathcal{S} :

$$F(z) = \lambda \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} (f_k(x_i^{\text{val}}) + z - y_i^{\text{val}})^2 + (1 - \lambda) \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} (f_k(x_i^{\text{val}}) + z - \hat{f}^{\text{init}}(x_i^{\text{val}}))^2$$

whose closed-form minimizer is

$$z^* = \lambda (\bar{y}_{\mathcal{S}}^{\text{val}} - \bar{f}_{k,\mathcal{S}}^{\text{val}}) + (1 - \lambda) (\bar{f}_{\mathcal{S}}^{\text{init}} - \bar{f}_{k,\mathcal{S}}^{\text{val}}) \quad (5)$$

where $\bar{y}_{\mathcal{S}}^{\text{val}}$, $\bar{f}_{k,\mathcal{S}}^{\text{val}}$, and $\bar{f}_{\mathcal{S}}^{\text{init}}$ denote the averages of y_i^{val} , $f_k(x_i^{\text{val}})$, and $\hat{f}^{\text{init}}(x_i^{\text{val}})$ over $i \in \mathcal{S}$, respectively. Here, z^* is a weighted correction that shifts the falsified model toward both the true labels and the ensemble’s predictive consensus.

A correction is accepted only if it reduces the Brier loss of the falsified model k on \mathcal{S} by at least δ . If accepted, we export the same mean shift z^* to the corresponding signed disagreement region on the test set. To define this region, we construct the signed disagreement sets on the test set analogously:

$$\mathcal{S}_{\text{test}}^> = \left\{ i \in \{1, \dots, n_{\text{test}}\} : f_m(x_i^{\text{test}}) - f_j(x_i^{\text{test}}) > \varepsilon \right\},$$

$$\mathcal{S}_{\text{test}}^< = \left\{ i \in \{1, \dots, n_{\text{test}}\} : f_j(x_i^{\text{test}}) - f_m(x_i^{\text{test}}) > \varepsilon \right\}.$$

When $\mathcal{S} = \mathcal{S}^>$ (resp. $\mathcal{S} = \mathcal{S}^<$) on validation, we export the shift to $\mathcal{S}^{\text{test}} = \mathcal{S}_{\text{test}}^>$ (resp. $\mathcal{S}^{\text{test}} = \mathcal{S}_{\text{test}}^<$).

For classification tasks, corrected predictions are clipped to $[0, 1]$ to preserve probabilistic validity. Iterations continue until no further verified corrections can be applied, yielding a reconciled ensemble with reduced internal disagreement and improved calibration consistency. Algorithm 3 summarizes the full procedure.

Experimental Setup

Datasets

We evaluate four high-stakes *classification* tasks (Table 1): Adult (Becker and Kohavi 1996), COMPAS (Larson et al. 2016), and two Folktables tasks, Folk Mobility and Folk Travel (Ding et al. 2021). All datasets use a 60/20/20 train/validation/test split and results are averaged over 10 random seeds.

Dataset	Size	Features	Target
Adult (UCI)	32,561	97	Income > 50K
COMPAS	6,172	351	2-year Recidivism
Folk Mobility	34,491	96	Mobility
Folk Travel	88,071	71	Travel Time

Table 1: Summary of classification datasets used in experiments. Feature counts reflect processed features, where categorical features were encoded as binary features.

Evaluation Metrics

We evaluate methods according to: (i) **predictive performance**, (ii) **predictive multiplicity**, and (iii) **neighborhood agreement**. Results are reported as mean \pm one standard deviation over 10 random seeds.

Predictive Performance. We report test-set metrics to ensure reconciliation does not degrade predictive quality: **Accuracy** (thresholding probabilities at 0.5) and the **Brier score** (mean squared error between predicted probabilities and binary labels). For each method, we compute these metrics using its ensemble prediction $\hat{f}(x^{\text{test}}) \in [0, 1]$ on $\mathcal{D}^{\text{test}}$.

Algorithm 3 Pairwise Reconciliation

Require: Model predictions $\{f_m(x_i^{\text{val}})\}_{m=1..M, i=1..n_{\text{val}}}$, $\{f_m(x_i^{\text{test}})\}_{m=1..M, i=1..n_{\text{test}}}$, labels $\{y_i^{\text{val}}\}_{i=1..n_{\text{val}}}$, disagreement threshold ε , batch size B , minimum region mass α , objective weight λ , acceptance tolerance δ , disagreement tolerance η , maximum iterations T_{max} .

- 1: Compute mean prediction on validation: $\hat{f}^{\text{init}}(x_i^{\text{val}}) \leftarrow \frac{1}{M} \sum_{m=1}^M f_m(x_i^{\text{val}})$ for all $i \in \{1, \dots, n_{\text{val}}\}$.
- 2: Initialize $t \leftarrow 0$.
- 3: **while** $t < T_{\text{max}}$ **do** \triangleright Compute pairwise disagreements on validation
 - 4: **for all** m and j such that $m < j$ **do**
 - 5: $D[m, j] \leftarrow \frac{1}{n_{\text{val}}} \sum_{i=1}^{n_{\text{val}}} |f_m(x_i^{\text{val}}) - f_j(x_i^{\text{val}})|$.
 - 6: **end for**
 - 7: **if** $\max_{m,j:m < j} D[m, j] < \eta$ **then**
 - 8: **break out of while loop**
 - 9: **end if**
 - 10: Select top- B disagreeing pairs \mathcal{P}_t according to $D[m, j]$.
 - 11: $\text{acceptedAtLeastOnce} \leftarrow \text{false}$.
 - 12: **for all** $(m, j) \in \mathcal{P}_t$ **do**
 - 13: $\mathcal{S}^> \leftarrow \{i : f_m(x_i^{\text{val}}) - f_j(x_i^{\text{val}}) > \varepsilon\}$.
 - 14: $\mathcal{S}^< \leftarrow \{i : f_j(x_i^{\text{val}}) - f_m(x_i^{\text{val}}) > \varepsilon\}$. \triangleright Define validation disagreement set \mathcal{S}
 - 15: **if** $|\mathcal{S}^>| \geq |\mathcal{S}^<|$ **then**
 - 16: $\mathcal{S} \leftarrow \mathcal{S}^>$ and $\sigma \leftarrow >$.
 - 17: **else**
 - 18: $\mathcal{S} \leftarrow \mathcal{S}^<$ and $\sigma \leftarrow <$.
 - 19: **end if**
 - 20: **if** $|\mathcal{S}| < \alpha$ **then**
 - 21: **skip to the next pair in** \mathcal{P}_t .
 - 22: **else** \triangleright Determine which model is falsified
 - 23: $k \leftarrow \arg \max_{u \in \{m, j\}} \text{BrierLoss}_{\mathcal{S}}(f_u)$.
 - 24: $\bar{f}_{k,\mathcal{S}}^{\text{val}} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} f_k(x_i^{\text{val}})$,
 - 25: $\bar{y}_{\mathcal{S}}^{\text{val}} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} y_i^{\text{val}}$,
 - 26: $\bar{f}_{\mathcal{S}}^{\text{init}} \leftarrow \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \hat{f}^{\text{init}}(x_i^{\text{val}})$.
 - 27: Compute closed-form correction z^* from Eq. (5).
 - 28: For all $i \in \mathcal{S}$, set $f'_k(x_i^{\text{val}}) \leftarrow \text{clip}_{[0,1]}(f_k(x_i^{\text{val}}) + z^*)$.
 - 29: **if** $\text{BrierLoss}_{\mathcal{S}}(f'_k) < \text{BrierLoss}_{\mathcal{S}}(f_k) - \delta$ **then** \triangleright Transfer correction to test set
 - 30: Accept: for all $i \in \mathcal{S}$, set $f_k(x_i^{\text{val}}) \leftarrow f'_k(x_i^{\text{val}})$.
 - 31: **if** $\sigma = >$ **then**
 - 32: $\mathcal{S}^{\text{test}} \leftarrow \{i : f_m(x_i^{\text{test}}) - f_j(x_i^{\text{test}}) > \varepsilon\}$.
 - 33: **else**
 - 34: $\mathcal{S}^{\text{test}} \leftarrow \{i : f_j(x_i^{\text{test}}) - f_m(x_i^{\text{test}}) > \varepsilon\}$.
 - 35: **end if**
 - 36: $\text{acceptedAtLeastOnce} \leftarrow \text{true}$.
 - 37: For all $i \in \mathcal{S}^{\text{test}}$, set
 - 38: $f_k(x_i^{\text{test}}) \leftarrow \text{clip}_{[0,1]}(f_k(x_i^{\text{test}}) + z^*)$.
 - 39: **end if**
 - 40: **end if**
 - 41: **end for**
 - 42: **if not** $\text{acceptedAtLeastOnce}$ **then**
 - 43: **break** \triangleright no accepted corrections; models already sufficiently agree
 - 44: **end if**
 - 45: $t \leftarrow t + 1$.
 - 46: **end while**
 - 47: **return** reconciled $\{f_m(x_i^{\text{val}})\}_{m=1..M, i=1..n_{\text{val}}}$ and reconciled $\{f_m(x_i^{\text{test}})\}_{m=1..M, i=1..n_{\text{test}}}$.

Predictive Multiplicity. We report four disagreement-based metrics computed from the set of model predictions $\{f_m(x^{\text{test}})\}_{m=1}^M$ on the test set. We define: (i) **Variance**, the mean over x^{test} of $\text{Var}_m[f_m(x^{\text{test}})]$; (ii) **Ambiguity**, the mean over x^{test} of $\max_m f_m(x^{\text{test}}) - \min_m f_m(x^{\text{test}})$; (iii) **Discrepancy**, the maximum over model pairs (m, j) of the mean absolute difference $\frac{1}{n_{\text{test}}} \sum_{x^{\text{test}} \in \mathcal{D}^{\text{test}}} |f_m(x^{\text{test}}) - f_j(x^{\text{test}})|$; and (iv) **Disagreement Rate**, the mean over model pairs (m, j) of the fraction of test points for which $|f_m(x^{\text{test}}) - f_j(x^{\text{test}})| > \varepsilon$, with $\varepsilon = 0.05$.

Neighborhood Agreement. We report **Local Conditional Absolute Error (LCAE)**, computed using validation neighborhoods, as a kNN-based proxy for local conditional error. We interpret lower LCAE as better **neighborhood agreement**, i.e., closer alignment between a test prediction and outcomes among nearby validation points. For each test point x_i^{test} , let $\mathcal{N}_{30}(x_i^{\text{test}}) \subset \mathcal{D}^{\text{val}}$ denote its $k = 30$ nearest neighbors in the validation set (in the preprocessed feature space). We define

$$\text{LCAE}_{30}(x_i^{\text{test}}) = \frac{1}{|\mathcal{N}_{30}(x_i^{\text{test}})|} \sum_{x_j^{\text{val}} \in \mathcal{N}_{30}(x_i^{\text{test}})} \left| \hat{f}(x_i^{\text{test}}) - y_j^{\text{val}} \right|, \quad (6)$$

where $\hat{f}(x_i^{\text{test}})$ is the method’s final ensemble prediction at x_i^{test} and $y_j^{\text{val}} \in \{0, 1\}$ is the validation label for neighbor x_j^{val} . We report the average over the test set:

$$\text{LCAE} = \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \text{LCAE}_{30}(x_i^{\text{test}}),$$

where lower values indicate better neighborhood agreement.

Experimental Protocol

To approximate a Rashomon set, we train a diverse pool of probabilistic classifiers in scikit-learn (logistic regression, random forests, extra trees, gradient boosting, and multi-layer perceptrons) using fixed hyperparameter grids. We evaluate candidates on the validation set and retain the top $M = 25$ models with the lowest validation Brier score.

Each dataset is split into 60%/20%/20% train/validation/test with stratification on the target label. The validation split is used (i) to define neighborhoods for neighborhood agreement (LCAE) and local patching, (ii) to select outlier-correction thresholds, (iii) to evaluate candidates for Rashomon set selection, and (iv) to identify disagreement regions and verify pairwise corrections. All results are reported as mean \pm one standard deviation over 10 random seeds.

Baselines

These baselines produce a single predictor from the Rashomon-set predictions $\{f_m(x)\}_{m=1}^M$:

- **Soft Voting** (ensemble mean): $\hat{f}(x) = \frac{1}{M} \sum_{m=1}^M f_m(x)$.
- **Majority Voting**: binarize each model at 0.5, average votes, and interpret the vote fraction as a probability.

- **Best Single Model**: choose the model in the Rashomon set with the highest validation accuracy (breaking ties by lower validation Brier score) and use it for all test points.
- **Random Selection**: for each test point, sample one model index uniformly at random and use its prediction.

Since these methods do not modify the underlying model predictions $\{f_m\}$, the ensemble-wide disagreement metrics (Var/Amb/Disc/Disag) are computed on the same underlying set $\{f_m\}$ and are thus identical across aggregation baselines.

Our approaches

Outlier Correction (OC). We evaluate OC, as in Algorithm 1. Outliers are identified using the Rashomon mean prediction and thresholds $(\tau_{\text{low}}, \tau_{\text{high}})$ that come from hyperparameter tuning. Training outlier labels are flipped and the ensemble is retrained; validation outlier labels are replaced with the Rashomon mean prediction.

Local Patching (LP). To evaluate LP, for each test point x^{test} , we form a neighborhood $N(x^{\text{test}}) \subset \mathcal{D}^{\text{val}}$ via k -nearest neighbors in the preprocessed feature space. Each model is locally corrected via the one-sided residual rule in Eq. (2), with threshold τ_{bias} , chosen via hyperparameter tuning on the validation set.

Pairwise Reconciliation (PR). To evaluate PR, we run Algorithm 3 on each trained Rashomon set, using the validation split to identify the top- B most-disagreeing pairs, form disagreement regions, and accept corrections only when they reduce validation Brier loss by at least δ . Hyperparameters $(\varepsilon, B, \alpha, \lambda, \delta, \eta, T_{\text{max}})$ are selected via validation tuning, and we report performance, multiplicity, and neighborhood agreement (LCAE) on the resulting reconciled test predictions.

Combinations. We evaluate combinations of our methods, specifically, OC+LP, OC+PR, PR+LP, and OC+PR+LP.

Results

Tables 2–5 summarize the relative performance of all methods across the four datasets in Table 1, reporting mean percent change ($\Delta\%$) relative to soft voting over 10 random seeds. Improvements are shown in green and degradations in pink.

We evaluate methods along three dimensions: predictive accuracy, neighborhood agreement (LCAE@30), and predictive multiplicity (Variance, Ambiguity, Discrepancy, and Disagreement Rate). The results highlight that predictive multiplicity is multi-dimensional.

Predictive Performance. Across all datasets, the combined methods preserve and often improve predictive accuracy relative to soft voting. In particular, PR+LP and OC+PR+LP achieve accuracy that is comparable to or better than PR alone in most settings. For example, OC+PR+LP improves accuracy on Adult, Folk Mobility, and Folk Travel while substantially reducing disagreement. On COMPAS,

accuracy changes remain small across all reconciliation-based variants, indicating that disagreement reduction does not come at the cost of substantial performance degradation.

Local Patching alone tends to incur modest accuracy losses, which is not intuitive and could potentially arise from overfitting based on our parameter choices. We used only 5 nearest neighbors to define the neighborhood, which is likely too small, and could be better chosen in future work. When LP is paired with reconciliation in our experiments, the losses are largely mitigated. Aggregation baselines such as majority voting and best single model show inconsistent accuracy and no systematic improvements.

Reduction of Predictive Multiplicity. Predictive multiplicity is most effectively reduced through reconciliation-based methods, but the results show that the most robust reductions arise from combinations rather than reconciliation alone. While PR alone drives dramatic reductions in disagreement (e.g., -99.67% on Adult and -99.77% on COMPAS), PR+LP and OC+PR+LP achieve similarly large reductions across all disagreement metrics while offering additional benefits in local agreement. PR+LP and OC+PR+LP achieve disagreement reductions comparable to PR alone, often exceeding -87% across datasets.

In contrast, Local Patching alone produces only moderate reductions in disagreement, and Outlier Correction without reconciliation does not consistently reduce ensemble-wide multiplicity. Aggregation baselines leave disagreement metrics unchanged, underscoring that modifying predictions rather than merely aggregating them substantially resolves predictive multiplicity.

Neighborhood Agreement (LCAE@30). Neighborhood agreement reveals the clearest advantage of combining methods. Local Patching is the primary driver of LCAE@30 improvements, consistently reducing neighborhood error across all datasets. However, LP alone does not eliminate disagreement and can negatively affect accuracy. When LP is combined with reconciliation, these limitations are addressed. PR+LP consistently achieves the largest reductions in LCAE@30 among reconciliation-based methods, such as -13.96% on Adult, -10.39% on Folk Mobility, and -6.00% on Folk Travel. OC+PR+LP also yields substantial improvements in neighborhood agreement while maintaining strong disagreement reduction and stable accuracy.

Pairwise reconciliation alone yields only modest LCAE improvements, indicating that reducing disagreement does not necessarily correct shared local bias. This distinction highlights the complementary roles of reconciliation and local patching: reconciliation aligns models with one another, while patching aligns predictions with outcomes.

Combined Effects and Practical Trade-offs. Taken together, the results show that no single mechanism is sufficient to address all aspects of predictive reliability. Pairwise reconciliation eliminates ensemble-wide disagreement, but it is insufficient for improving neighborhood agreement on its own. Local patching corrects local bias, and outlier correction stabilizes downstream corrections by improving label quality.

The strongest overall trade-offs are consistently obtained by PR+LP and OC+PR+LP, achieving low disagreement,

Method	Δ Acc \uparrow	Δ LCAE \downarrow	Δ Var \downarrow	Δ Amb \downarrow	Δ Disc \downarrow	Δ Disag \downarrow
Soft Voting	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%
Random Selection	-0.95%	+0.04%	+0.00%	+0.00%	+0.00%	+0.00%
Majority Voting	+0.07%	-24.15%	+0.00%	+0.00%	+0.00%	+0.00%
Best Single Model	+0.37%	-7.17%	+0.00%	+0.00%	+0.00%	+0.00%
LP	-0.49%	-13.88%	-16.98%	-27.80%	-33.73%	-39.95%
OC + LP	+0.00%	-12.08%	-7.55%	-17.45%	-24.34%	-27.49%
PR	+0.28%	-5.08%	-98.11%	-84.59%	-90.14%	-99.67%
OC + PR	+0.32%	-5.00%	-83.02%	-55.78%	-70.96%	-67.20%
PR + LP	-0.35%	-13.96%	-96.23%	-86.91%	-92.44%	-98.66%
OC + PR + LP	+0.25%	-11.46%	-98.11%	-87.52%	-92.68%	-98.77%

Table 2: Adult dataset.

Method	Δ Acc \uparrow	Δ LCAE \downarrow	Δ Var \downarrow	Δ Amb \downarrow	Δ Disc \downarrow	Δ Disag \downarrow
Soft Voting	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%
Random Selection	-3.29%	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%
Majority Voting	+0.12%	-15.12%	+0.00%	+0.00%	+0.00%	+0.00%
Best Single Model	+0.43%	-2.81%	+0.00%	+0.00%	+0.00%	+0.00%
LP	-1.82%	-7.96%	-3.06%	-12.04%	-18.65%	-18.93%
OC + LP	+0.10%	-4.99%	+5.10%	+0.00%	-2.25%	-4.95%
PR	-0.73%	-2.11%	-98.98%	-89.62%	-93.69%	-99.77%
OC + PR	+0.26%	-2.13%	-82.65%	-56.28%	-68.02%	-57.10%
PR + LP	-2.27%	-7.18%	-96.94%	-88.31%	-92.59%	-98.46%
OC + PR + LP	-0.74%	-4.48%	-97.96%	-88.83%	-92.87%	-99.03%

Table 3: COMPAS dataset.

Method	Δ Acc \uparrow	Δ LCAE \downarrow	Δ Var \downarrow	Δ Amb \downarrow	Δ Disc \downarrow	Δ Disag \downarrow
Soft Voting	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%
Random Selection	-1.17%	-0.03%	+0.00%	+0.00%	+0.00%	+0.00%
Majority Voting	-0.76%	-28.82%	+0.00%	+0.00%	+0.00%	+0.00%
Best Single Model	-0.06%	-1.19%	+0.00%	+0.00%	+0.00%	+0.00%
LP	-0.80%	-10.06%	+5.26%	-10.40%	-12.25%	-20.82%
OC + LP	-0.07%	-8.04%	+7.89%	-3.34%	-6.76%	-10.18%
PR	+0.99%	-1.28%	-92.11%	-69.67%	-78.71%	-90.53%
OC + PR	+0.88%	-2.48%	-71.05%	-48.78%	-63.00%	-57.67%
PR + LP	-0.08%	-10.39%	-81.58%	-69.91%	-79.58%	-89.26%
OC + PR + LP	+0.91%	-7.85%	-86.84%	-68.91%	-78.32%	-87.69%

Table 4: Folk Mobility dataset.

Method	Δ Acc \uparrow	Δ LCAE \downarrow	Δ Var \downarrow	Δ Amb \downarrow	Δ Disc \downarrow	Δ Disag \downarrow
Soft Voting	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%
Random Selection	-4.16%	+0.00%	+0.00%	+0.00%	+0.00%	+0.00%
Majority Voting	-0.54%	-7.55%	+0.00%	+0.00%	+0.00%	+0.00%
Best Single Model	+2.61%	-1.99%	+0.00%	+0.00%	+0.00%	+0.00%
LP	-4.19%	-6.15%	+8.93%	-11.07%	-13.57%	-21.86%
OC + LP	-0.14%	-2.28%	+5.36%	-0.77%	-1.76%	-5.04%
PR	+2.58%	-1.62%	-94.64%	-76.93%	-81.07%	-96.06%
OC + PR	+2.36%	-1.45%	-80.36%	-54.75%	-63.39%	-57.08%
PR + LP	-2.34%	-6.00%	-85.71%	-75.78%	-81.29%	-93.94%
OC + PR + LP	+2.42%	-2.85%	-92.86%	-76.62%	-81.07%	-95.59%

Table 5: Folk Travel dataset.

improved neighborhood agreement, and competitive accuracy.

References

- Becker, B.; and Kohavi, R. 1996. Adult. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5XW20>.
- Black, E.; Raghavan, M.; and Solon, B. 2022. Model multiplicity: Opportunities, concerns, and solutions. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 850–863.
- Breiman, L. 2001. Statistical modeling: The two cultures. *Statistical Science*, 16(3): 199–231.
- Breunig, M. M.; Kriegel, H.-P.; Ng, R. T.; and Sander, J. 2000. LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 93–104.
- Coker, B.; Rudin, C.; and King, G. 2021. A Theory of Statistical Inference for Ensuring the Robustness of Scientific Results. *Management Science*, 67(10).
- D’Amour, A.; Heller, K.; Moldovan, D.; Adlam, B.; Alipanahi, B.; Beutel, A.; Chen, C.; Deaton, J.; Eisenstein, J.; Hoffman, M. D.; et al. 2022. Underspecification presents challenges for credibility in modern machine learning. *Journal of Machine Learning Research*, 23(226): 1–61.
- Ding, F.; Hardt, M.; Miller, J.; and Schmidt, L. 2021. Retiring Adult: New Datasets for Fair Machine Learning. In *Advances in Neural Information Processing Systems*, volume 34, 6478–6490.
- Du, A. Y.; Ngo, D. D.; and Wu, Z. S. 2024. Reconciling model multiplicity for downstream decision making. arXiv preprint. arXiv:2405.19667.
- Fisher, A.; Rudin, C.; and Dominici, F. 2019. All Models are Wrong, but Many are Useful: Learning a Variable’s Importance by Studying an Entire Class of Prediction Models Simultaneously. *Journal of Machine Learning Research*, 20(177): 1–81.
- Kleinberg, J.; Ludwig, J.; Mullainathan, S.; and Rambachan, A. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences*, 118(22): e2018340118.
- Kobylińska, K.; Krzyżiński, M.; Machowicz, R.; Adamek, M.; and Biecek, P. 2024. Exploration of the Rashomon set assists trustworthy explanations for medical data. *IEEE Journal of Biomedical and Health Informatics*, 28(11): 6454–6465.
- Larson, J.; Angwin, J.; Kirchner, L.; and Mattu, S. 2016. How We Analyzed the COMPAS Recidivism Algorithm. *ProPublica*.
- Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2012. Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1): 1–39.
- Long, C.; et al. 2024. Individual Arbitrariness and Group Fairness. In *Advances in Neural Information Processing Systems*, volume 36.
- Marx, C.; Calmon, F.; Ustun, B.; and Ustun, B. 2020. Predictive multiplicity in classification. In *Proceedings of the International Conference on Machine Learning*, 6765–6774. PMLR.
- McCullagh, P.; and Nelder, J. A. 1989. *Generalized Linear Models, Second Edition*. CRC Press.
- Roth, A.; Tolbert, A.; Weinstein, A., Scott; and Weinstein, S. 2023. Reconciling individual probability forecasts. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 101–110.
- Rudin, C. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5): 206–215.
- Rudin, C.; Zhong, C.; Semenova, L.; Seltzer, M.; Parr, R.; Liu, J.; Katta, S.; Donnelly, J.; Chen, H.; and Boner, Z. 2024. Amazing Things Come From Having Many Good Models. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Semenova, L.; Rudin, C.; and Parr, R. 2022. On the Existence of Simpler Machine Learning Models. In *ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)*.
- Watson-Daniels, J.; Parkes, D. C.; and Ustun, B. 2023. Predictive Multiplicity in Probabilistic Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 10306–10314.
- Xin, R.; Zhong, C.; Chen, Z.; Takagi, T.; Seltzer, M.; and Rudin, C. 2022. Exploring the whole Rashomon set of sparse decision trees. *Advances in Neural Information Processing Systems*, 35: 14071–14084.
- Zhao, Y.; Nasrullah, Z.; and Li, Z. 2019. Pyod: A python toolbox for scalable outlier detection. *Journal of Machine Learning Research*, 20(96): 1–7.