Conformal Data Contamination Tests for In-distribution Data Acquisition

Anonymous authors
Paper under double-blind review

Abstract

The amount of quality data in many machine learning tasks is limited to what is available locally to data owners. The set of quality data can be expanded through trading or sharing with external data agents. However, external data may be contaminated or introduce undesirable sample diversity which can degrade performance of personalized machine learning tasks, as in diagnosis of a rare disease or recommendation systems. Therefore, data buyers need quality guarantees prior to data acquisition. Previous works primarily rely on distributional assumptions about data from different agents, relegating quality checks to post-hoc steps involving costly data valuation procedures. We propose a distribution-free, contamination-aware data acquisition framework that, by inspecting only a small volume of data, identifies external data agents whose data is most valuable for model personalization. To achieve this, we introduce novel two-sample testing procedures, preceding full data acquisition, grounded in rigorous theoretical foundations for conformal outlier detection, to determine whether an agent's data exceeds a contamination threshold. The proposed tests, termed *conformal data contamination tests*, remain valid under arbitrary contamination levels and the novel Storey-type test provably enables finite-sample false discovery rate control via the Benjamini-Hochberg procedure. Empirical evaluations across diverse collaborative learning scenarios demonstrate the robustness and effectiveness of our approach. Overall, the conformal data contamination test distinguishes itself as a generic procedure for aggregating data with statistically rigorous quality guarantees.

1 Introduction

In many real-world machine learning applications the amount of quality data for training is limited to what is locally available to the data owners. Collaborative learning techniques have shown some potential by allowing multiple participants to contribute data and jointly train a model in a privacy-preserving manner (McMahan et al., 2017), however, relying solely on distributed and diverse private data from collaborators in the learning process significantly affects *personalization* (Blum et al., 2017). Unlike federated learning (McMahan et al., 2017), several practical scenarios consider a data owner having specific learning goals, and not being interested in having a *well-generalized* model that works for the rest of the participants. An example would be a small, energy-efficient model that is specialized for a certain type of task and there is no interest to perform well on other tasks (Belcak et al., 2025). This indicates that *not all* data, but only quality data, meaning data with *specific attributes*, are relevant for the learning agent, and motivates this work to resolve challenges of quality data acquisition.

We distinguish data agents with distribution mismatch into three categories: (i) honest data agents with unintentionally biased data, (ii) dishonest data agents selling artificial data, and (iii) malicious data agents generating adversarial data (Szegedy et al., 2014; Zhang et al., 2023). For the first case, an example is a hospital interested in training a model to detect a rare disease based on an X-ray image of a patient. Since the disease is rare, the hospital may not have sufficient local data to train an accurate prediction model, but when buying data from other hospitals, or a *data marketplace* (Fernandez, 2023), it is important to ensure similar data conditions, as different X-ray machines and softwares may have been used to take the measurements (Al-qaness et al., 2024). An example of the second case is fake social media accounts sold

on a data marketplace as genuine users. For the third case, an example is recommendation systems with fake review injection (Nawara et al., 2025). These examples outline the need for having principled data acquisition/sharing protocols with quality guarantees.

Quality data acquisition approaches have motivated novel data trading or sharing mechanisms, such as *data markets*, where the data agent can prioritize access to relevant data for better personalization (Fernandez et al., 2020). In a nutshell, we assume that *in-distribution* implies relevant data, similar to the training samples, for improving personalization. In contrast, *out-of-distribution* (OOD) indicates outliers and unseen samples, including data with potential contamination, which may hold value for better generalization. A participant's benefit from the data of other participants hinges on careful data selection per its learning goal, as outlined in the earlier real-world examples, leading to the development of efficient data valuation techniques, such as Shapley value (Ghorbani & Zou, 2019) and its variants, amongst others. Unfortunately, these approaches are still limited: first, they are model-specific and computationally expensive; second, they are often used *post-hoc*, i.e., after collecting the data, which might be contaminated and of less value; and third, detecting the optimal set of collaboration partners without making any distributional assumptions on their data is challenging. This raises some fundamental questions in developing data sharing mechanisms prior to training. Furthermore, in a networked system, any of such mechanisms should be compatible with emerging peer-to-peer collaboration protocols between decentralized data sources and AI-powered tools, such as *model context protocol* (MCP) (Hou et al., 2025). From the perspective of MCP, the data agents are the local data sources and an MCP host then wants to access the data from the local data sources which are relevant for personalization. In this premise, the focus of this work is to develop an agile data trading framework between data agents that identifies valuable training data, *without making any distributional assumption*, through the formalization of guarantees to data contamination tests.

Data contamination model: Consider that a local data agent observes data $\mathcal{D}_0 = \{(X_i, Y_i)\}_{i=1}^n$ of n input-output pairs, where $X_i \in \mathbb{R}^d$ is the observed feature and Y_i is the response with $Z_i = (X_i, Y_i) \sim P_0$ for an unknown local distribution P_0. In the data sharing platform, there are K other data agents with each their own local data $\mathcal{D}_k = \{(X_i^k, Y_i^k)\}_{i=1}^{m_k}$ of m_k input-output pairs with $Z_i^k = (X_i^k, Y_i^k) \sim \tilde{P}_k$ for unknown local distributions $\tilde{P}_k$. Each of the unknown local distributions can be decomposed as $\tilde{P}_k = (1 - \pi_k)P_0 + \pi_k P_k$ where P_k is a proper outlier distribution and $\pi_k \in [0, 1]$ is the contamination factor (Blanchard et al., 2010). If the learning goal is better personalization, as is the focus of our work, finding the best data agent(s) for data acquisition naturally boils down to learning who has the least contaminated data.

General data sharing procedure and challenges: *First*, upon the availability of data from K data agents in the data sharing platform, the data requester, say the data agent 0, must decide from which data agent(s) to acquire more data to improve personalization (Lee et al., 2018). Once the first problem of who to pair with is resolved, the data sharing procedure is executed in rounds, where, in each round, a batch of data is acquired and used to decide the data acquisition policy in the next round; the process ends in either of the following cases: (i) when the data has been exhausted, (ii) the data purchasing budget has been depleted, or (iii) the data agent 0 decides that acquiring more data is not necessary (beneficial). *Second*, after the termination of data sharing, data agent 0 uses a data subset selection technique, for instance a complex data valuation technique (Koh & Liang, 2017; Ghorbani & Zou, 2019; Yoon et al., 2020), or a simpler outlier detection method (Schölkopf et al., 2001; Hawkins et al., 2002; Ruff et al., 2018), to filter out OOD datapoints and consequently improve personalization.

Based on the data contamination model, the first challenge can be posed as testing null hypotheses $H_0^1, \ldots, H_0^K$ where $H_0^k : \pi_k \leq \pi_{\text{th}}$ for a user-specified contamination factor threshold $\pi_{\text{th}} \in [0, 1)$. The idea is that we reject H_0^k when we have sufficient evidence that the data from the k-th data agent is contaminated beyond the threshold we allow. A theoretical (but impractical) solution to this problem has been considered in Blanchard et al. (2010), and some existing works are occupied with estimating the contamination factor (Ramaswamy et al., 2016; Perini et al., 2023); however, Ramaswamy et al. (2016) provide no distributional guarantees, and Perini et al. (2023) uses a complex Bayesian approach that requires distributional and prior assumptions.

Main Contributions: The objective of this work is to improve the local model performance through quality data acquisition during training. This is done by employing a data selection strategy that identifies agents for collaboration through novel data contamination tests, based on low-complexity outlier detection algorithms.

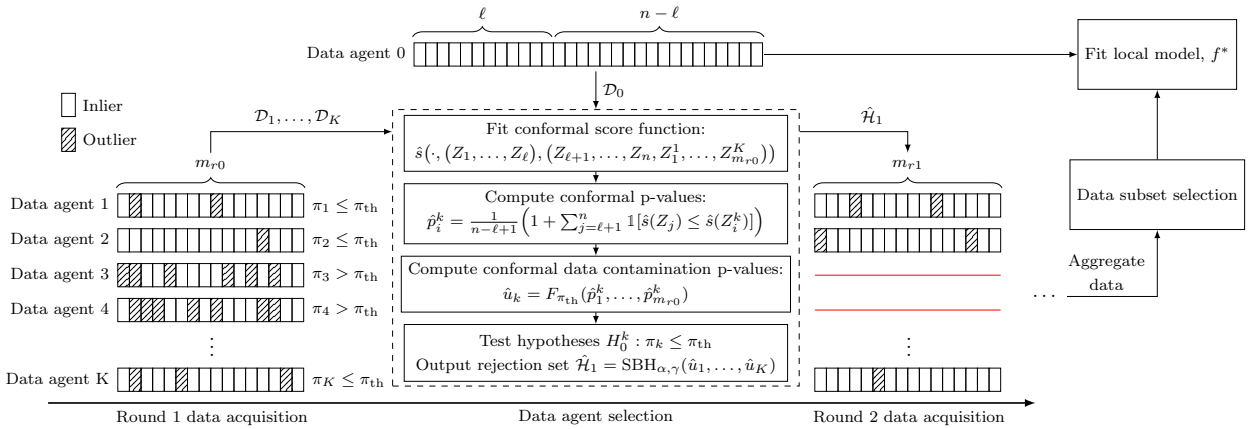

Figure 1: *Proposed data acquisition procedure:* In the first round, m_{r0} samples are received at data agent 0 from each of the other K data agents. Then, a conformal score function $\hat{s}$ is fitted to the data with a split of the local data into ℓ training data and $n-\ell$ calibration data, conformal p-values $\hat{p}_i^k$ are computed for the acquired data, the proposed conformal data contamination p-values $\hat{u}_k$ are evaluated as functions $F_{\pi_{\text{th}}}$ of the conformal p-values, and these data contamination p-values are used for testing the null hypotheses $H_0^k : \pi_k \leq \pi_{\text{th}}$ with Storey's Benjamini-Hochberg ($\text{SBH}_{\alpha,\gamma}$) procedure, resulting in a data agent rejection set $\hat{\mathcal{H}}_1$. In the following round, data agent 0 only acquires data from those not in the rejection set. Finally, a data subset selection technique is used to filter away OOD data, and the local model is trained.

The key contribution of this technique is its practicality and low-complexity relative to expensive data valuation techniques. Additionally, the proposed procedure allows for flexible and interpretable adjustment of the trade-off between the amount of data acquired and the relevance of the acquired data for the personalized learning goal. Specifically, we propose a distribution-free solution for testing $H_0^1, \ldots, H_0^K$ with false discovery rate (FDR) control guarantees that builds on the ideas of conformal outlier detection (Bates et al., 2023; Marandon et al., 2024), and apply the developed tools to the data acquisition scenario, as outlined in Figure 1. The main contributions are summarized as follows:

- We propose novel two-sample testing procedures based on four types of conformal p-value combination test statistics. This enables us to define four types of conformal data contamination p-values for testing $H_0^k : \pi_k \leq \pi_{\text{th}}$, $k = 1, \ldots, K$, without any distributional assumptions which provably controls the false rejection probability, and generalize the combination tests of Bates et al. (2023) from the special case of $\pi_{\text{th}} = 0$ to any $\pi_{\text{th}} \in [0, 1)$. Moreover, the tests provide a lot of flexibility with the choice of the conformal score, and as such are compatible with a wide range of low-complexity outlier detection methods. This contribution has independent interest and can potentially have practical significance for many applications. In this work, we highlight how these tests can be used to improve personalization through in-distribution data acquisition in data contamination scenarios.
- We show that the proposed p-value sequence, based on a Storey-type combination test statistic, for testing $H_0^1, \ldots, H_0^K$ is positive regression dependent on a subset (PRDS) thereby allowing for FDR control using the Benjamini-Hochberg (BH) procedure.
- We propose a data sharing procedure developed for personalization that uses the conformal data contamination tests to initialize collaboration of data agents through data sharing while providing theoretical guarantees.
- Numerical experiments are conducted on medical image datasets, as well as a label noise scenario on the MNIST dataset, to validate the effectiveness of the conformal data contamination tests and the proposed data sharing procedure.

We reiterate the fact that the setup is fundamentally different from the one of federated learning: we are not concerned with privacy-aware distributed model training that would work for all involved parties, but rather acquisition of quality data that contributes to personalized learning goals.

Other related works: Several works towards collaborative machine learning ecosystems focus on efficient mechanism design that enables data sharing amongst heterogeneous data agents. Based on rich literature from economics and game theory, novel value attribution frameworks have been investigated to answer fundamental questions on data collection and sharing, which precede compute-intensive model training. For instance, Huang et al. (2023) discussed interactions between data agents when designing mechanisms for data collection in collaborative learning. Studies on *data marketplaces* started evaluating context and specific goals of data agents to develop data acquisition processes, such as in Lu et al. (2024), where a buyer's objective is fulfilled through collection of statistically relevant data for training a predictive model, and in Capitaine et al. (2024), where the total welfare is maximized through collaboration. While *pricing* on data value has been a key incentive for assuring quality data acquisition in most works, Ananthakrishnan et al. (2024) extended it to understand how much test data to use for verification, and how to tune hyperparameters once data collection is done for offering machine learning services.

A line of work in conformal inference, that is related to but different from ours, deals with *corruption* in the calibration data, which in the general sense means that the calibration data is not exactly distributed according to the null distribution P_0. Such *corruption* violates the fundamental exchangeability assumption that conformal inference relies on, however, approximately valid inference is still possible. The methodologies in this line of work is relevant for downstream uncertainty quantification following (imperfect) in-distribution data acquisition. Tibshirani et al. (2019) proposed weighted conformal prediction and applied it to handle covariate shift. Subsequently, weighted conformal prediction was considered for label shift in Podkopaev & Ramdas (2021) and for generalized covariate shift and posterior drift in Wang & Qiao (2025). A detailed theoretical analysis of conformal prediction beyond exchangeability was conducted in Barber et al. (2023). Recently *robust* conformal prediction has been considered: Bashari et al. (2025) proposed a technique to avoid *corruption* by detecting outliers in the calibration data, and Feldman & Romano (2024) considered scenarios with access to privileged information (not available at test time) that explains the distribution shift.

2 Preliminaries

Notations : Let $[n] = \{1, \ldots, n\}$ and $[n]_0 = \{0, \ldots, n\}$ denote index sets for $n \geq 1$, and in an abuse of notation let $[n]/m = \{1/m, \ldots, n/m\}$ and $m + [n] = \{m + 1, \ldots, m + n\}$.

2.1 Multiple testing

Consider a multiple testing scenario in which K hypotheses $H_0^1, \ldots, H_0^K$ are tested, and denote by $p_1, \ldots, p_K$ the p-values. We define S as the number of rejected null hypotheses that are actually false, and V as the number of rejected null hypotheses that are actually true (type I error). The total number of rejected hypotheses is given by $R = V + S$, while K_0 represents the total number of true null hypotheses.

A widely studied error rate in multiple hypothesis testing is the FDR, introduced by Benjamini & Hochberg (1995). It is defined as $\text{FDR} = \mathbb{E}[V/\max(1, R)]$. Similarly, the power of a test, also known as the true discovery rate (TDR), is defined as $\mathbb{E}[S/\max(1, K - K_0)]$. The most well-known method for controlling the FDR is the BH procedure. Let $p_{(1)} \leq \cdots \leq p_{(K)}$ denote the ordered p-values. The BH procedure at level $q^* \in (0, 1)$ rejects hypotheses with indices $j \leq \kappa$ where $\kappa = \max\{j \in [K] : p_{(j)} \leq q^* j/K\}$. Equivalently, the rejection set is $\hat{\mathcal{H}}_1 = \{j \in [K] : p_j \leq p_{(\kappa)}\}$. When the p-values are PRDS (see Section A.1), the BH procedure controls the FDR at level $q^* K_0/K$ (Benjamini & Yekutieli, 2001). There exists adaptive procedures which estimate K_0, defined as $\hat{K}_0$, and accordingly adjusts the procedure by running it at level $q^* = \alpha K/\hat{K}_0$, for significance level $\alpha \in (0, 1)$ (Storey et al., 2003; Benjamini et al., 2006). With Storey's BH procedure, $\kappa_{\alpha,\gamma} = \max\{j \in [K] : p_{(j)} \leq \alpha j/\hat{K}_0\}$, where $\hat{K}_0/K = \sum_{j=1}^{K} \mathbb{1}[p_j > \gamma]/(1 - \gamma)$ with the hyperparameter $\gamma \in (0, 1)$ controlling the bias-variance trade-off, and we denote the rejection set of this procedure by $\text{SBH}_{\alpha,\gamma}(p_1, \ldots, p_K) = \{j \in [K] : p_j \leq p_{(\kappa_{\alpha,\gamma})}\}$.

2.2 Conformal outlier detection

In the conformal outlier detection setting, a number of test points $Z_{n+1}, \ldots, Z_{n+m}$, as well as a null sample $Z_1, \ldots, Z_n$ are observed. It is assumed that the null sample datapoints are exchangeable and from an

unknown distribution P_0. Then, a real-valued conformal score function, $\hat{s}$, is defined satisfying the following permutation invariance property

$$\hat{s}(\cdot) := \hat{s}(\cdot, (Z_1, \ldots, Z_\ell), (Z_{\sigma(\ell+1)}, \ldots, Z_{\sigma(n+m)})) = \hat{s}(\cdot, (Z_1, \ldots, Z_\ell), (Z_{\ell+1}, \ldots, Z_{n+m})),$$

for any permutation σ of $\{\ell+1, \ldots, n+m\}$, and $0 \leq \ell < n$ (Marandon et al., 2024). Large scores $\hat{s}_{n+i} = \hat{s}(Z_{n+i})$ indicate conformity with P_0; small scores indicate outliers. As in Marandon et al. (2024), we assume that the inliers in the test sample are exchangeable with the calibration data $Z_{\ell+1}, \ldots, Z_n$, conditional on the outliers. For $i \in [m]$, the conformal p-value associated to the test point Z_{n+i} is

$$\hat{p}_i = \frac{1}{n-\ell+1}\Big(1 + \sum_{j=\ell+1}^{n} \mathbb{1}[\hat{s}_j \leq \hat{s}_{n+i}]\Big). \tag{1}$$

Assuming that the conformal scores are continuously distributed and that Z_{n+i} is an inlier independent of $Z_1, \ldots, Z_\ell$, then $\hat{p}_i$ is uniformly distributed on $[n-\ell+1]/(n-\ell+1)$, and one has that $\lfloor \beta(n-\ell+1) \rfloor/(n-\ell+1) \leq \mathbb{P}(\hat{p}_i \leq \beta) \leq \beta$, for $\beta \in (0,1)$. This shows that $\hat{p}_i$ is a *marginally* valid p-value with almost exact control of the false rejection probability (Marandon et al., 2024).

3 Main results

In the first round of data sharing as in Figure 1, the data from the k-th agent yields conformal p-values $\hat{p}_1^k, \ldots, \hat{p}_m^k$ (setting $m_{r0} = m$ for notational convenience). In the following, we consider a single sequence of conformal p-values, and so to simplify notations we denote this sequence by $\hat{p}_1, \ldots, \hat{p}_m$. We are interested in testing the null hypothesis $H_0 \ : \ \pi \leq \pi_{\text{th}}$ using this sequence of conformal p-values. Such tests can be referred to as combination tests as they use a test statistic which combines a sequence of p-values. We also refer to these tests as two-sample tests since the conformal p-values, and in turn the combination test statistics, are computed using two samples of data, i.e., the null sample and the test sample.

We denote by $\mathrm{B}^m_{\pi_{\text{th}}}(k)$ the probability mass function of a binomial distribution evaluated at k with parameters π_{th} and m, by $F_{\text{NHG}}(x; n-\ell+k, n-\ell, k-r)$ the cumulative distribution function (CDF) of the negative hypergeometric distribution evaluated at $x \in [n-\ell]_0$ with population size $n-\ell+k$, number of success states $n-\ell$, and number of failures $k-r$, by $F_{\chi^2_{2k}}$ the CDF of the chi-square distribution with $2k$ degrees of freedom, and by F_{IH_k} the CDF of the Irwin-Hall distribution with parameter k. Further, let $\bar{\mathcal{H}}_1 = \{i \in [m] : Z_{n+i} \not\sim P_0\}$ denote the set of outliers. We now state one of our main results, defining valid p-values for two natural test statistics.

Theorem 3.1. *Let $H_0 : \pi \leq \pi_{\text{th}}$, and let $\hat{p}_i$ for $i \in [m]$ denote the conformal p-values as introduced in Section 2.2. The following are valid p-values for H_0:*

$$\hat{u}^{\text{storey}} = \sum_{k=0}^{T^{\text{storey}}} \mathrm{B}^m_{\pi_{\text{th}}}(k) + \sum_{k=T^{\text{storey}}+1}^{m} \mathrm{B}^m_{\pi_{\text{th}}}(k) F_{\text{NHG}}\big(\lfloor \lambda(n-\ell+1) \rfloor - 1; n-\ell+k, n-\ell, k-T^{\text{storey}}\big), \tag{2}$$

where $T^{\text{storey}} = \sum_{i=1}^m \mathbb{1}[\hat{p}_i > \lambda]$ and $\lambda \in [n-\ell]/(n-\ell+1)$ is a hyperparameter, moreover, if $\mathbb{P}(\hat{p}_i \leq \lambda, \forall i \in \bar{\mathcal{H}}_1) = 1$ and $\pi = \pi_{\text{th}}$ the p-value is exact;

$$\hat{u}^{\text{quantile}} = \sum_{k=0}^{i_0} \mathrm{B}^m_{\pi_{\text{th}}}(k) + \sum_{k=i_0+1}^{m} \mathrm{B}^m_{\pi_{\text{th}}}(k) F_{\text{NHG}}\big(T^{\text{quantile}} - 1; n-\ell+k, n-\ell, k-i_0\big), \tag{3}$$

where $T^{\text{quantile}} = (n-\ell+1)\hat{p}_{(m-i_0)}$ with $\hat{p}_{(m-i_0)}$ denoting the $(m-i_0)$-th smallest conformal p-value with hyperparameter $i_0 \in [m-1]_0$, moreover, if $\mathbb{P}(\hat{p}_i \leq \hat{p}_{(m-i_0)}, \forall i \in \bar{\mathcal{H}}_1) = 1$ and $\pi = \pi_{\text{th}}$ the p-value is exact.

Proof. The proof is given in Section A of the appendix. □

Remark 3.2. The strategy of the proof is to use the law of total probability to decompose the rejection probability under the null into a sum over the number of inliers. Subsequently, an inequality is made by discarding contributions from the conformal p-values of outliers. Finally, using the marginal distribution of the order statistics of the conformal p-values, which specifically is a negative hypergeometric distribution, see Gazin et al. (2024) and Biscio et al. (2025), allows for an explicit expression for the rejection probability under the null, which in turn defines a p-value.

The test statistic T^{storey} is motivated by the Storey estimator of π by Storey (2002), while T^{quantile} is motivated by the quantile estimator of π by Benjamini et al. (2006).

More general test statistics are also allowed, however, constructing valid p-values require knowledge of the distribution of the test statistic under the null. In the following result, we present two other test statistics and their associated p-values which are asymptotically valid.

Theorem 3.3. *Let $H_0 : \pi \leq \pi_{\text{th}}$ and let $\hat{p}_i$ for $i \in [m]$ denote the conformal p-values as introduced in Section 2.2. Then, asymptotically as $n - \ell \to \infty$ the following are valid p-values for H_0:*

$$\hat{u}^{\text{fisher}} = \pi_{\text{th}}^m + \sum_{k=1}^{m} \mathrm{B}_{\pi_{\text{th}}}^m(k) F_{\chi^2_{2k}}\left(2k + \frac{-T^{\text{fisher}} + 2k\big(\log(n-\ell+1) - 1\big)}{\sqrt{1 + k/(n-\ell)}}\right), \tag{4}$$

where $T^{\text{fisher}} = -2\sum_{i=1}^{m}\left(\log\left(\frac{1}{n-\ell+1}\right) - \log(\hat{p}_i)\right)$*;*

$$\hat{u}^{\text{sum}} = \pi_{\text{th}}^m + \sum_{k=1}^{m} \mathrm{B}_{\pi_{\text{th}}}^m(k) F_{\mathrm{IH}_k}\left(\frac{T^{\text{sum}} + k(\sqrt{1 + k/(n-\ell)} - 1)/2}{\sqrt{1 + k/(n-\ell)}}\right), \tag{5}$$

where $T^{\text{sum}} = \sum_{i=1}^{m} \hat{p}_i$*.*

Proof. The proof is found in Section A of the appendix. □

Remark 3.4. The strategy of the proof is similar as that of Theorem 3.1, however, since the distribution of T^{fisher} or T^{sum} where $\hat{p}_i$ are all inliers is not explicitly known, we employ the asymptotic result Theorem S1 of Bates et al. (2023). The general theorem is actually more general than presented here, as one can consider a general class of test statistics of the form $\sum_{i=1}^{m} G(\hat{p}_i)$ requiring that $G(U)$ has finite moments for $U \sim \mathrm{Unif}([0,1])$ (Bates et al., 2023). The general statement with precise conditions required on G is given in the appendix.

The test statistic T^{fisher} is motivated by Fisher's combination test (Fisher, 1925), and T^{sum} is another classical combination test statistic (Vovk & Wang, 2020).

With the p-values of Eqs. (2)-(5), tests for H_0 can be conducted without requiring any assumptions on the null distribution, P_0, and outlier distribution, P_1. Effective conformal data contamination testing relies on conformal scores that clearly separate inliers from outliers, leading to small conformal p-values for outliers and enabling exact control when $\pi = \pi_{\text{th}}$, as seen in the details in the appendix. The key parameter is ℓ, which balances test power: larger ℓ improves score quality, while larger $n - \ell$ enhances distribution approximation. Finally, having a large m in general also improves the power of the conformal data contamination tests as more evidence against the null can be aggregated. This is opposite to the case of conformal outlier detection for which Mary & Roquain (2022) studied how larger m results in a loss of power.

In case multiple test sets are observed, we have multiple sets of conformal p-values, each of which gives a conformal data contamination p-value. We consider here how to do multiple testing with this sequence of conformal data contamination p-values while proving FDR control guarantees. Motivated by the result of Bates et al. (2023) that conformal p-values are PRDS, we have shown that the Storey conformal data contamination p-values are indeed also PRDS. Hence, we can simply apply the BH procedure on the conformal data contamination p-values to guarantee FDR control.

Theorem 3.5. *Let $\hat{u}_1^{\text{storey}}, \ldots, \hat{u}_K^{\text{storey}}$ be K Storey conformal data contamination p-values as in Eq. (2), with the k-th derived from conformal p-values $\hat{p}_1^k, \ldots, \hat{p}_m^k$ assuming independence of the K test sets $\mathcal{D}_k = \{(X_i^k, Y_i^k)\}_{i=1}^m$. Then, the Storey conformal data contamination p-values are PRDS.*

Proof. The proof is found in Section A of the appendix. □

Remark 3.6. While Theorem 3.5 only refers to the Storey conformal data contamination p-values, we expect the result can be generalized to a broad class of conformal data contamination p-values, herein the ones presented in this paper, and we leave this as a line of future investigation.

Selecting the test: Based on theoretical considerations and numerical experiments, we give the following recommendations: For small calibration sets (e.g. $n-\ell < 100$), the Storey and Quantile statistics are generally preferred as they offer finite-sample guarantees; however, both require the selection of a hyperparameter, which can be challenging in practice. Since the Storey and Quantile statistics can take m and $n-\ell$ distinct values, respectively, we recommend the Storey statistic when $m > n-\ell$ and vice versa. When ample calibration data are available, the Fisher or Sum statistics are recommended as they are hyperparameter free. The Fisher statistic is suited for low contamination thresholds, whereas the Sum statistic is appropriate for higher contamination thresholds.

4 Proposed data sharing procedure

In this section, we describe the proposed data sharing procedure in further details and discuss its complexity and limitations.

Taking the perspective of the 0-th data agent, we consider a scenario where in the first round each data agent $k \in [K]$ shares a test sample $\{Z_i^k\}_{i=1}^{m_{r0}}$. We assume for simplicity that the number of shared datapoints m_{r0} is the same for all data agents, however this assumption is not required for the proposed methodology. We model data coming from the k-th data agent by the contamination model (Blanchard et al., 2010):

$$Z_i^k \sim (1-\pi_k)P_0 + \pi_k P_k$$

where $\pi_k \in [0,1]$ is the contamination factor.[1] As in Section 2.2, we require that the calibration data and the inliers in the test sample are exchangeable, conditional on the outliers, for each $k \in [K]$. Outliers $Z_i^k \not\sim P_0$ are not of interest for the 0-th data agent as we focus on improving personalization, here, by selecting relevant collaborators.

We proceed by testing null hypotheses $H_0^k : \pi_k \leq \pi_{\text{th}}$ given a threshold π_{th} on how much contamination we tolerate. To do this, we begin by computing conformal p-values as in Eq. (1)

$$\hat{p}_i^k = \frac{1}{n-\ell+1}\Big(1 + \sum_{j=\ell+1}^{n} \mathbb{1}[\hat{s}_j \leq \hat{s}_i^k]\Big),$$

where $\hat{s}_i^k = \hat{s}(Z_i^k)$ are the test scores. The conformal p-values $\hat{p}_1^k, \ldots, \hat{p}_{m_{r0}}^k$ are then mapped to conformal data contamination p-values $\hat{u}_k$, as presented in Section 3. Finally, for a specified significance level, $\alpha \in (0,1)$, the hypotheses are tested using an adaptive BH procedure, e.g., Storey's BH procedure described in Section 2.1, yielding a set of non-rejected data agents, $\hat{\mathcal{H}}_0 = [K] \setminus \text{SBH}_{\alpha,\gamma}(\hat{u}_1, \ldots, \hat{u}_K)$ for Storey's hyperparameter $\gamma \in (0,1)$. The statistical guarantee, provided by the proposed tests, ensures that data agents observing a distribution highly relevant for the personalized learning goal, are rarely neglected in the collaboration.

After determining which data agents to collaborate with in the following round(s), the 0-th data agent once again acquires data $Z_{m_{r0}+i}^k$ for $i \in [m_{r1}]$ and $k \in \hat{\mathcal{H}}_0$, and can then aggregate all the received data, followed by data subset selection for instance using techniques from data valuation (Koh & Liang, 2017; Ghorbani & Zou, 2019; Yoon et al., 2020). To remain lightweight, we employ conformal outlier detection for this part, specifically, when fitting the local prediction model, we do not use acquired data that is detected as outliers using Storey's BH procedure with significance level β and Storey hyperparameter ζ, i.e., the set $\text{SBH}_{\beta,\zeta}(\{\hat{p}_i^k : i \in [m_{r0}], k \in [K]\} \cup \{\hat{p}_i^k : i \in m_{r0} + [m_{r1}], k \in \hat{\mathcal{H}}_0\})$. The selected data is combined with the n local datapoints, and the 0-th data agent trains its local model. The proposed procedure is visualized in Figure 1 and also summarized in Procedure 1. An overview of the scenario variables and hyperparameters is given in Section B, and in Section D.5 a data-driven approach to hyperparameter selection is outlined.

[1] We make the assumption here that P_k are proper novelty distributions with respect to P_0: there exists no decomposition of the form $P_k = (1-\zeta)Q + \zeta P_0$ where Q is a probability distribution and $0 < \zeta \leq 1$, see Blanchard et al. (2010); Zhu et al. (2023) for details.

Procedure 1

Input: local data $\{Z_i\}_{i=1}^{\ell}$ and $\{Z_i\}_{i=\ell+1}^{n}$, conformal score method $\hat{s}$, model class f, incoming data per round per data agent m_{r0} and m_{r1}, contamination threshold π_{th}, significance level α, conformal non-contamination statistic T.
Output: optimized local model, f^*.
Receive data (round 1) from other data agents $\{Z_i^k\}_{i=1}^{m_{r0}}$, $k \in [K]$.
Learn conformal score function $\hat{s}(\cdot) = \hat{s}(\cdot, (Z_1, \ldots, Z_\ell), (Z_{\ell+1}, \ldots, Z_n, Z_1^1, \ldots, Z_{m_{r0}}^K))$.
Compute calibration scores $\hat{s}_{\ell+1}, \ldots, \hat{s}_n$ where $\hat{s}_i = \hat{s}(Z_i)$.
For each $k \in [K]$, compute test scores, $\hat{s}_{n+1}^k, \ldots, \hat{s}_{n+m_{r0}}^k$ where $\hat{s}_{n+j}^k = \hat{s}(Z_j^k)$.
For each $k \in [K]$, compute conformal p-values, $\hat{p}_1^k, \ldots, \hat{p}_{m_{r0}}^k$.
Evaluate conformal data contamination p-values, $\hat{u}_k$ (see Section 3).
Collaborate in the following round with data agents in $\hat{\mathcal{H}}_0 = [K] \setminus \text{SBH}_{\alpha,\gamma}(\hat{u}_1, \ldots, \hat{u}_K)$.
Receive data (round 2) from other data agents $\{Z_{m_{r0}+i}^k\}_{i=1}^{m_{r1}}$, $k \in \hat{\mathcal{H}}_0$.
Run data subset selection on all the received data, specifically, discard data $\text{SBH}_{\beta,\zeta}(\{\hat{p}_i^k : i \in [m_{r0}], k \in [K]\} \cup \{\hat{p}_i^k : i \in m_{r0} + [m_{r1}], k \in \hat{\mathcal{H}}_0\})$.
Use all local data $\{Z_i\}_{i=1}^{n}$ together with the selected data to train the model, yielding f^*.

Complexity: One of the strengths of the proposed data acquisition protocol is its low computational complexity due to preceding the expensive model training procedure. Note that computing a conformal data contamination p-value as defined in Eqs. (2)-(5) requires at most m evaluations of known CDFs: for Quantile and Storey the CDF of a negative hypergeometric distribution, for Fisher the CDF of a chi-square distribution, and for Sum the CDF of an Irwin-Hall distribution. The most computationally complex part of the data acquisition protocol is fitting of the conformal score function $\hat{s}$ but notably the methodology complies with a wide range of conformal scores with the choice open to the practitioner. Further, the procedure scales well to a large number of data agents by virtue of the Benjamini-Hochberg procedure which was designed for large scale testing.

Limitations: The proposed data acquisition procedure has some limitations. First, the conformal data contamination tests require exchangeability of calibration and test data to be theoretically valid, restricting use in settings like time series. Second, we chose to focus on the conformal data contamination tests, rather than aspects of data subset selection, however, this could be explored further in other works. Lastly, the proposed data sharing procedure is not privacy-aware, limiting applicability to scenarios with strict privacy constraints. However, a slightly altered procedure, which replaces the first step of data sharing by federated learning of the conformal score function, can avoid privacy-leakage for data agents not selected for subsequent collaboration.

5 Numerical experiments and analysis

In this section, we numerically evaluate the proposed methodology on two data contamination scenarios. In the first scenario, retinal fundus images are used to classify severity of diabetic retinopathy and the inliers are drawn from RetinaMNIST (Liu et al., 2022) with contaminated data drawn from EyePACS (Dugas et al., 2015; Gulshan et al., 2016). In the second scenario, images of handwritten digits are used to classify the digits (1, 4, and 7) and the inliers are drawn noise-free from MNIST (Lecun et al., 1998) with contaminated data having incorrect labels. First we introduce the data setup, and present the baseline approaches. Then, the conformal data contamination tests are analyzed numerically. Finally, the performance of the proposed data sharing procedure is evaluated on the classification tasks.

5.1 Data setup and baselines

The RetinaMNIST and EyePACS datasets consists of 1600 and 35126 retinal fundus images, respectively. Each image in RetinaMNIST is $28 \times 28 \times 3$ with each pixel taking values in $[255]_0$ and has an associated label $y \in [4]_0$ indicating the severity of diabetic retinopathy ranging from no diabetic retinopathy to proliferate

retinopathy. The EyePACS data has varying resolutions, so we crop and downsample the images to the $28 \times 28 \times 3$ resolution.

The MNIST dataset restricted to digits 1, 4, and 7 consists of 21994 images of handwritten digits. Each image is 28×28 with each pixel taking values in $[255]_0$ and has an associated label $y \in [2]_0$ indicating the digit.

We normalize the images to the interval $[0, 1]$ by dividing pixel-wise with 255. For all the numerical experiments we set $n = 100$, $\ell = 60$, $m = m_{r0} = m_{r1} = 40$, $K = 10$, $\lambda = \lfloor (n-\ell)/8 \rfloor / (n-\ell+1)$, $i_0 = \lfloor m/3 \rfloor$, and $\gamma = 0.5$. Unless otherwise specified, we simulate the data contamination factors independent and identically distributed (iid) according to $\pi_k \sim \text{Uniform}([0, 1])$. In all cases we sample randomly n in-distribution datapoints for the 0-th data agent. Meanwhile, the k-th data agent observes $2m$ datapoints and the number of OOD datapoints is sampled from $\text{B}^{2m}_{\pi_k}$.

As a conformal score, we will use the *positive unlabelled* (PU) classification approach advocated in Marandon et al. (2024). For simplicity, the PU classifier is chosen as logistic regression (LR), and we use the standard implementation in *scikit-learn* with balanced class weights. We define the score to be class-wise in the sense that

$$\hat{s}((X, Y), (Z_1, \ldots, Z_\ell), (Z_{\ell+1}, \ldots, Z_n, Z^1_1, \ldots, Z^K_m)) = \sum_{i \in \mathcal{Y}} \mathbb{1}[Y = i] \; \hat{s}_i(X, X^{Y=i}_{\text{train}}, X^{Y=i}_{\text{cal+test}}),$$

where $\mathcal{Y}$ denotes the set of labels, $X^{Y=i}_{\text{train}} = (X_j : Y_j = i, j \in [\ell])$ denotes the subset of $X_1, \ldots, X_\ell$ for which $Y_j = i$ for $j \in [\ell]$, and likewise $X^{Y=i}_{\text{cal+test}} = (X_{\ell+j} : Y_{\ell+j} = i, j \in [n-\ell]) \cup (X^k_j : Y^k_j = i, j \in [m], k \in [K])$. Specifically, $\hat{s}_i(X, X^{Y=i}_{\text{train}}, X^{Y=i}_{\text{cal+test}})$ denotes the predicted probability that (X, i) has label 1, i.e., is an inlier, using LR fitted to $X^{Y=i}_{\text{train}}$ with labels 1 and $X^{Y=i}_{\text{cal+test}}$ with labels -1.

As a classifier we consider support vector classification (SVC) (Boser et al., 1992; Cortes & Vapnik, 1995) for simplicity and robustness, again using the standard implementation in *scikit-learn* (Pedregosa et al., 2011).

We consider the following baselines: (i) *Oracle 1:* The complete oracle knows the true contamination factors and exactly which datapoints are contaminated. Accordingly, it selects for collaboration in the second round the data agents with data contamination factor π_k below a specified threshold π_{th}, and then uses only the in-distribution data for subsequent model training. (ii) *Oracle 2:* The partial oracle knows the true contamination factors but not which datapoints are contaminated. Accordingly, it selects for collaboration in the second round the data agents with data contamination factor π_k below a specified threshold π_{th}, and then uses all collected data for subsequent model training. (iii) *Random:* The naïve baseline randomly selects K_{budget} agents to collaborate with in the second round, and then uses all collected data for subsequent model training. The proposed method is benchmarked either without data subset selection (*Storey*/*Quantile*/*Fisher*/*Sum*), or using conformal outlier detection with Storey's BH procedure (*Storey**/*Quantile**/*Fisher**/*Sum**), setting the significance level to $\beta = 0.1$ and Storey's hyperparameter to $\zeta = 0.5$.

In Section C of the appendix we provide an ablation study with synthetic Gaussian data. We also provide additional numerical experiments with a variety of conformal scores for the data contamination scenario presented here, as well as seven other data contamination scenarios, in Section D, herein, also validating scalability to larger dataset sizes in Section D.4.

5.2 Conformal data contamination tests

In this section the conformal data contamination tests are numerically analyzed on the retinal fundus images. We showcase in Figure 2 the (empirical) CDF of the conformal data contamination p-values with $\pi_{\text{th}} = 0.1$ and for $\pi = 0.1, 0.3$. From Figure 2a we observe that the conformal data contamination tests are superuniform under the null, as dictated by the theory in Section 3, and notice that the Quantile p-value is nearly standard uniform, meanwhile the Fisher p-value is the most conservative in this case. Figure 2b shows that in this case the Storey test is the most powerful followed by the Sum and Fisher tests, and finally the weakest is the Quantile test.

We report in Table 1(a) the (empirical) TDR for the conformal data contamination with varying data contamination thresholds $\pi_{\text{th}} \in \{0, 0.1, 0.2\}$ at significance level $\alpha = 0.1$ when using Storey's BH procedure with Storey's hyperparameter $\gamma = 0.5$. We observe the highest TDR with the Storey test and a generally low

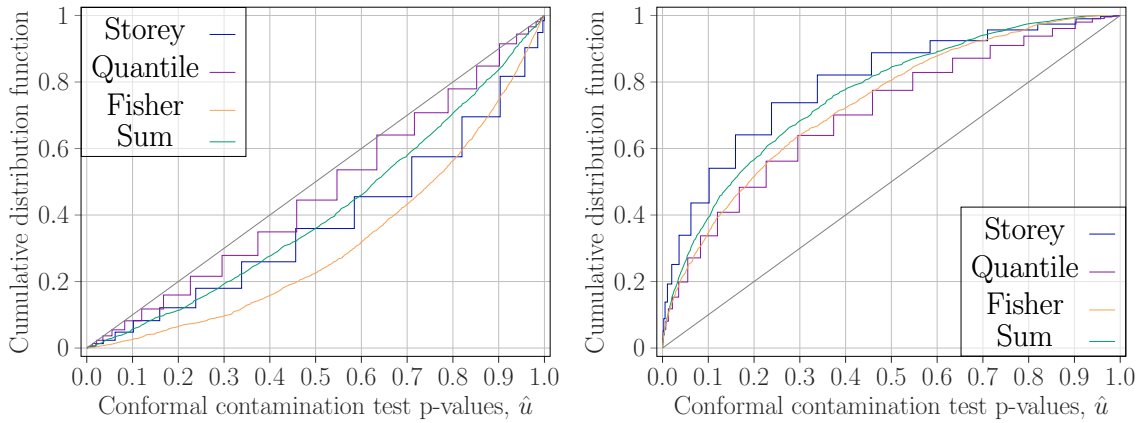

(a) $\pi = 0.1$. (b) $\pi = 0.3$.

Figure 2: Empirical CDF of conformal data contamination p-values with $\pi_{\text{th}} = 0.1$ for (a) $\pi = 0.1$ and (b) $\pi = 0.3$. The gray line is the CDF of a standard uniform.

Table 1: TDR and FDR computed from 2000 simulations for varying data contamination thresholds, π_{th}, at significance level $\alpha = 0.1$ using Storey's BH procedure with parameter $\gamma = 0.5$. (a) TDR with data contamination factor simulated as $\pi_k \sim \text{Uniform}([0, 1])$, and (b) FDR where for $K/2$ of the data agents $\pi_k = \pi_{\text{th}}$ while for the others $\pi_k = 1 - \pi_{\text{th}}$.

		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$			$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(a) TDR	Storey	0.9128	0.8147	0.7033	(b) FDR	Storey	0.1064	0.0708	0.0489
	Quantile	0.8104	0.7851	0.7537		Quantile	0.1120	0.1074	0.1050
	Fisher	0.8890	0.7576	0.6077		Fisher	0.0641	0.0281	0.0153
	Sum	0.8778	0.7994	0.7142		Sum	0.0890	0.0658	0.0458

TDR with the Quantile test, and notably, the Storey and Fisher tests perform well at low π_{th} but tend to decrease rapidly in TDR as the data contamination threshold increases, while the decrease is less pronounced with the Quantile and Sum tests.

In Table 1(b) we numerically validate the FDR control provided by the conformal data contamination tests on the boundary of the null hypothesis. Here, $K_0 = K/2$ data agents have data contamination factor $\pi_k = \pi_{\text{th}}$, and the remaining data agents have $\pi_k = 1 - \pi_{\text{th}}$ so approximately half of the test data is contaminated. Through this construction, we are on the boundary of the null hypothesis, i.e., there are no $\pi_k < \pi_{\text{th}}$ contributing to excess conservativeness from being in the interior of the null hypothesis. Once again, we used Storey's BH procedure at significance level $\alpha = 0.1$ and with Storey's hyperparameter $\gamma = 0.5$. The results show that the tests effectively controls the FDR at the nominal level, and we notice that the Quantile test is the least conservative among the four tests.

5.3 Collaborative data sharing

We show in Figure 3 the relation between the average number of collaborating data agents, $\mathbb{E}[|\hat{\mathcal{H}}_0|]$, and the average classification accuracy across 500 data simulations for both data contamination scenarios, with various hyperparameter settings $(\pi_{\text{th}}, \alpha) \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7\} \times \{0.05, 0.2, 0.5, 0.7\}$. The proposed data sharing procedures without data subset selection outperform no data sharing and the *random* baseline, even reaching the performance of *oracle 2*, showing the efficiency of the proposed methodology, and highlighting the accuracy improvements achieveable by carefully selecting data agents for collaboration.

The biggest gains in accuracy without data subset selection can be observed in Figure 3b showing the case of label noise on the MNIST data. Specifically, at $\mathbb{E}[|\hat{\mathcal{H}}_0|] \approx 6$ we achieve an increase in accuracy of up to 3 percentage points with the proposed methods compared to the *random* baseline. For the other scenario, at $\mathbb{E}[|\hat{\mathcal{H}}_0|] \approx 6$ we achieve an increase in accuracy just below 1 percentage point.

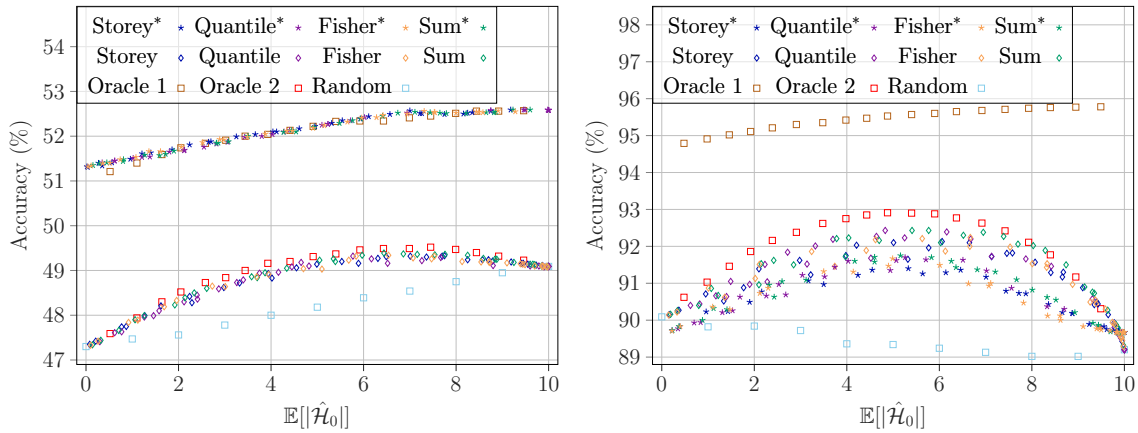

(a) Retinal fundus images. Accuracy without data sharing: 47.77 %.

(b) Handwritten digits. Accuracy without data sharing: 90.75 %.

Figure 3: Plot of the classification accuracy against the average number of collaborating data agents. Results with seven other datasets are available in Section D.

When including data subset selection via the lightweight conformal outlier detection, for the retinal fundus image scenario, the performance of the proposal reaches even that of *oracle 1*, however, the same is not observed for the label noise scenario, since the outlier detection problem proved more difficult for that data scenario. This shows that if the conformal score function is sufficiently powerful, conformal outlier detection can practically be used as a lightweight data subset selection technique.

6 Conclusion

We have presented a data sharing framework which selects collaboration partners using novel distribution-free testing procedures, named *conformal data contamination tests*. Leveraging conformal p-values, our method detects agents whose data exceed a contamination threshold while providing false discovery rate guarantees. This enables strategic acquisition of high-quality and personalized data while avoiding irrelevant or harmful data sources. Experiments validate the effectiveness and practicality of our approach.

Our framework opens up several promising directions for future work, including integration with incentive mechanisms in real-world data markets, and extension to continual learning with online data acquisition policies.

Broader Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Mohammed A. A. Al-qaness, Jie Zhu, Dalal AL-Alimi, Abdelghani Dahou, Saeed Hamood Alsamhi, Mohamed Abd Elaziz, and Ahmed A. Ewees. Chest X-ray images for lung disease detection using deep learning techniques: A comprehensive survey. *Archives of Computational Methods in Engineering*, 31(6):3267–3301, 2024. doi: 10.1007/s11831-024-10081-y.

Nivasini Ananthakrishnan, Stephen Bates, Michael I. Jordan, and Nika Haghtalab. Delegating data collection in decentralized machine learning. In *International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, pp. 478–486. PMLR, 2024.

Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816 – 845, 2023. doi: 10.1214/23-AOS2276.

Meshi Bashari, Matteo Sesia, and Yaniv Romano. Robust conformal outlier detection under contaminated reference data. In *International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 3091–3141. PMLR, 2025.

Stephen Bates, Emmanuel Candès, Lihua Lei, Yaniv Romano, and Matteo Sesia. Testing for outliers with conformal p-values. *The Annals of Statistics*, 51(1):149 – 178, 2023. doi: 10.1214/22-AOS2244.

Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. Small language models are the future of agentic AI, 2025. URL `https://arxiv.org/abs/2506.02153`.

Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300, 1995. ISSN 00359246.

Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4):1165 – 1188, 2001. doi: 10.1214/aos/1013699998.

Yoav Benjamini, Abba M. Krieger, and Daniel Yekutieli. Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507, 09 2006. ISSN 0006-3444. doi: 10.1093/biomet/93.3.491.

Christophe A. N. Biscio, Adrien Mazoyer, and Martin V. Vejling. Conformal novelty detection for replicate point patterns with FDR or FWER control. *Spatial Statistics*, 69, 2025. ISSN 2211-6753. doi: 10.1016/j.spasta.2025.100924.

Gilles Blanchard, Gyemin Lee, and Clayton Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11(99):2973–3009, 2010.

Avrim Blum, Nika Haghtalab, Ariel D Procaccia, and Mingda Qiao. Collaborative PAC learning. In *Advances in Neural Information Processing Systems*, volume 30 of *Proceedings of Machine Learning Research*. PMLR, 2017.

Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, COLT '92, pp. 144–152. Association for Computing Machinery, 1992. ISBN 089791497X. doi: 10.1145/130385.130401.

Aymeric Capitaine, Etienne Boursier, Antoine Scheid, Eric Moulines, Michael Jordan, El-Mahdi El-Mhamdi, and Alain Oliviero Durmus. Unravelling in collaborative learning. In *Advances in Neural Information Processing Systems*, Proceedings of Machine Learning Research. PMLR, 2024.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995. doi: 10.1007/BF00994018.

Emma Dugas, Jared, Jorge, and Will Cukierski. Diabetic retinopathy detection, 2015. URL `https://kaggle.com/competitions/diabetic-retinopathy-detection`.

Shai Feldman and Yaniv Romano. Robust conformal prediction using privileged information. In *Advances in Neural Information Processing Systems*, Proceedings of Machine Learning Research. PMLR, 2024.

Raul Castro Fernandez. Data-sharing markets: Model, protocol, and algorithms to incentivize the formation of data-sharing consortia. *Proc. ACM Manag. Data*, 1(2), 06 2023. doi: 10.1145/3589317.

Raul Castro Fernandez, Pranav Subramaniam, and Michael J. Franklin. Data market platforms: trading data assets to solve data problems. *Proc. VLDB Endow.*, 13(12):1933–1947, 7 2020. ISSN 2150-8097. doi: 10.14778/3407790.3407800.

Ronald A. Fisher. *Statistical Method for Research Workers*. Oliver & Boyd, Edinburgh, 1925.

Ulysse Gazin, Gilles Blanchard, and Etienne Roquain. Transductive conformal inference with adaptive scores. In *International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pp. 1504–1512. PMLR, 05 2024.

Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2242–2251. PMLR, 6 2019.

Varun Gulshan, Lily Peng, Marc Coram, Martin C. Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ramasamy Kim, Rajiv Raman, Philip C. Nelson, Jessica L. Mega, and Dale R. Webster. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Journal of American Medical Association*, 316(22):2402–2410, 12 2016. ISSN 0098-7484. doi: 10.1001/jama.2016.17216.

Simon Hawkins, Hongxing He, Graham Williams, and Rohan Baxter. Outlier detection using replicator neural networks. In *Data Warehousing and Knowledge Discovery*, pp. 170–180. Springer Berlin Heidelberg, 2002. ISBN 978-3-540-46145-6.

Xinyi Hou, Yanjie Zhao, Shenao Wang, and Haoyu Wang. Model context protocol (MCP): Landscape, security threats, and future research directions, 2025. URL `https://arxiv.org/abs/2503.23278`.

Baihe Huang, Sai Praneeth Karimireddy, and Michael Jordan. Evaluating and incentivizing diverse data contributions in collaborative learning. In *International Conference on Machine Learning Workshop: Federated Learning and Analytics in Practice: Algorithms, Systems, Applications, and Opportunities*, Proceedings of Machine Learning Research. PMLR, 2023.

Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1885–1894. PMLR, 2017.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, volume 31 of *Proceedings of Machine Learning Research*. PMLR, 2018.

Ruhan Liu, Xiangning Wang, Qiang Wu, Ling Dai, Xi Fang, Tao Yan, Jaemin Son, Shiqi Tang, Jiang Li, Zijian Gao, Adrian Galdran, J.M. Poorneshwaran, Hao Liu, Jie Wang, Yerui Chen, Prasanna Porwal, Gavin Siew Wei Tan, Xiaokang Yang, Chao Dai, Haitao Song, Mingang Chen, Huating Li, Weiping Jia, Dinggang Shen, Bin Sheng, and Ping Zhang. DeepDRiD: Diabetic retinopathy – grading and image quality estimation challenge. *Patterns*, 3(6):100512, 2022. ISSN 2666-3899. doi: https://doi.org/10.1016/j.patter.2022.100512.

Charles Lu, Baihe Huang, Sai Praneeth Karimireddy, Praneeth Vepakomma, Michael Jordan, and Ramesh Raskar. Data acquisition via experimental design for data markets. In *Advances in Neural Information Processing Systems*, Proceedings of Machine Learning Research. PMLR, 2024.

Ariane Marandon, Lihua Lei, David Mary, and Etienne Roquain. Adaptive novelty detection with false discovery rate guarantee. *The Annals of Statistics*, 52(1):157 – 183, 2024. doi: 10.1214/23-AOS2338.

David Mary and Etienne Roquain. Semi-supervised multiple testing. *Electronic Journal of Statistics*, 16(2): 4926 – 4981, 2022. doi: 10.1214/22-EJS2050.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 1273–1282. PMLR, 2017.

Dina Nawara, Ahmed Aly, and Rasha Kashef. Shilling attacks and fake reviews injection: Principles, models, and datasets. *IEEE Trans. Comput. Soc. Syst.*, 12(1):362–375, 2025. doi: 10.1109/TCSS.2024.3465008.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Lorenzo Perini, Paul-Christian Bürkner, and Arto Klami. Estimating the contamination factor's distribution in unsupervised anomaly detection. In *International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 27668–27679. PMLR, 2023.

Aleksandr Podkopaev and Aaditya Ramdas. Distribution-free uncertainty quantification for classification under label shift. In *Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pp. 844–853. PMLR, 2021.

Harish Ramaswamy, Clayton Scott, and Ambuj Tewari. Mixture proportion estimation via kernel embeddings of distributions. In *International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2052–2060. PMLR, 6 2016.

Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4393–4402. PMLR, 6 2018.

Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001. doi: 10.1162/089976601750264965.

John D. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(3):479–498, 08 2002. ISSN 1369-7412. doi: 10.1111/1467-9868.00346.

John D. Storey, Jonathan E. Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(1):187–205, 12 2003. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2004.00439.x.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, Proceedings of Machine Learning Research. PMLR, 2014.

Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. In *Advances in Neural Information Processing Systems*, volume 32 of *Proceedings of Machine Learning Research*. PMLR, 2019.

Vladimir Vovk and Ruodu Wang. Combining p-values via averaging. *Biometrika*, 107(4):791–808, 06 2020. ISSN 0006-3444. doi: 10.1093/biomet/asaa027.

Baozhen Wang and Xingye Qiao. Conformal prediction under generalized covariate shift with posterior drift. In *International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 4888–4896. PMLR, 2025.

Jinsung Yoon, Sercan Arik, and Tomas Pfister. Data valuation using reinforcement learning. In *International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 10842–10851. PMLR, 07 2020.

Shuhai Zhang, Feng Liu, Jiahao Yang, Yifan Yang, Changsheng Li, Bo Han, and Mingkui Tan. Detecting adversarial data by probing multiple perturbations using expected perturbation score. In *International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 41429–41451. PMLR, 2023.

Yilun Zhu, Aaron Fjeldsted, Darren Holland, George Landon, Azaree Lintereur, and Clayton Scott. Mixture proportion estimation beyond irreducibility. In *International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 42962–42982. PMLR, 2023.

A Conformal data contamination tests: theoretical details and proofs

We consider the scenario where we have observed a null sample $\mathcal{D}_{\text{null}} = \{Z_i\}_{i=1}^n$, $Z_i \sim P_0$, as well as a test sample $\mathcal{D}_{\text{test}} = \{Z_{n+j}\}_{j=1}^m$, $Z_{n+j} \sim P_j$. We call data from P_0 inliers and data from some other distribution P_1 outliers. We define $\bar{\mathcal{H}}_0 = \{j \in [m] : Z_{n+j} \sim P_0\}$, and $\bar{\mathcal{H}}_1 = [m] \setminus \bar{\mathcal{H}}_0$. Following the conformal outlier detection procedure as described in Section 2.2, conformal p-values are computed $\hat{p}_1, \ldots, \hat{p}_m$. As in Section 2.2 we shall assume that the conformal score is continuously distributed, or almost surely has no ties, and that the inliers are exchangeable conditioned on the outliers (Bates et al., 2023).

Two classical estimators of the contamination factor π are the Storey estimator of Storey et al. (2003) and the quantile estimator of Benjamini et al. (2006) defined respectively as

$$\hat{\pi}^{\text{storey}} = 1 - \frac{\sum_{i=1}^m \mathbb{1}[\hat{p}_i > \lambda]}{m(1-\lambda)},$$
$$\hat{\pi}^{\text{quantile}} = 1 - \frac{i_0 + 1}{m(1 - \hat{p}_{(m-i_0)})},$$

where $\hat{p}_{(1)} \leq \cdots \leq \hat{p}_{(m)}$ are the ordered conformal p-values. Here $\lambda \in [n-\ell]/(n-\ell+1)$ and $i_0 \in [m-1]_0$ are hyperparameters controlling the bias-variance trade-off. For these two estimators, the parts related to the data are $\sum_{i=1}^m \mathbb{1}[\hat{p}_i > \lambda]$ and $\hat{p}_{(m-i_0)}$, respectively, and so in the following we will use these as test statistics, noting that small values of these tests statistics will yield large estimates of π and so small values are evidence against the null hypothesis $H_0 : \pi \leq \pi_{\text{th}}$.

In the following, we give a detailed proof of Theorem 3.1.

Proof of Theorem 3.1. Since the test statistic T^{storey} is left-sided, its exact p-value is $\mathbb{P}_{H_0}(T^{\text{storey}} \leq t^{\text{storey}})$ where t^{storey} denotes the observed value of the random variable T^{storey}. By the law of total probability

$$\begin{aligned}\mathbb{P}_{H_0}(T^{\text{storey}} \leq t^{\text{storey}}) &= \sum_{k=0}^m \mathbb{P}_{H_0}(k)\mathbb{P}\Big(\sum_{i=1}^k \mathbb{1}[\hat{p}_i > \lambda] + \sum_{i=k+1}^m \mathbb{1}[\hat{p}_i > \lambda] \leq t^{\text{storey}} \mid \bar{\mathcal{H}}_0 = [k]\Big) \\ &\leq \sup_{\pi_0 \in [0,\pi_{\text{th}}]} \sum_{k=0}^m \mathrm{B}_{\pi_0}^m(k)\mathbb{P}\Big(\sum_{i=1}^k \mathbb{1}[\hat{p}_i > \lambda] \leq t^{\text{storey}} \mid \bar{\mathcal{H}}_0 = [k]\Big). \end{aligned} \tag{A.1}$$

The inequality is tight when $\hat{p}_i \leq \lambda$ for all $i \in \mathcal{H}_1$ which with a well chosen λ can occur if the conformal score separates well the inliers from the outliers. The distribution of the null data points is binomial with success probability $\pi_0 \in [0, \pi_{\text{th}}]$,

$$\mathrm{B}_{\pi_0}^m(k) = \binom{m}{k}(1-\pi_0)^k \pi_0^{m-k}.$$

The probability of the sum can be expressed in terms of the marginal CDF of the ordered conformal p-values as

$$\mathbb{P}\Big(\sum_{i=1}^k \mathbb{1}[\hat{p}_i > \lambda] \leq t^{\text{storey}} \mid \bar{\mathcal{H}}_0 = [k]\Big) = \begin{cases} 1 & \text{if } k \leq t^{\text{storey}}, \\ \mathbb{P}\big(\hat{p}_{(k-t^{\text{storey}})} \leq \lambda\big) & \text{otherwise,} \end{cases}$$

where $\hat{p}_{(1)} \leq \hat{p}_{(2)} \leq \cdots \leq \hat{p}_{(k)}$ are the ordered conformal p-values among the k p-values from the null. Notably, $\hat{p}_{(k-t^{\text{storey}})}$ is distributed according to a negative hypergeometric distribution with population size $n-\ell+k$, number of success states $n-\ell$, and number of failures $k-r$ supported on $\{1, 2, \ldots, n-\ell+1\}$, see Corollary A.1 of Biscio et al. (2025), and the CDF can be expressed as

$$\mathbb{P}\big(\hat{p}_{(k-t^{\text{storey}})} \leq \lambda\big) = F_{\text{NHG}}\big(\lfloor \lambda(n-\ell+1) \rfloor - 1; n-\ell+k, n-\ell, k-t^{\text{storey}}\big).$$

Finally, observe that the supremum in Eq. (A.1) occurs at π_{th}, since $\sum_{k=0}^l \mathrm{B}_{\pi_0}^m(k)$, $l \in \{0, \ldots, m\}$ is a non-decreasing function in π_0, and $\mathbb{P}(\sum_{i=1}^k \mathbb{1}[\hat{p}_i > \lambda] \leq t^{\text{storey}} \mid \bar{\mathcal{H}}_0 = [k])$ is a decreasing function of k as the indicator function is non-negative. This concludes the derivation of the p-value $\hat{u}^{\text{storey}}$

The proof of the validity of $\hat{u}^{\text{quantile}}$ follows immediately by noticing that for $x \in [n-\ell+1]/(n-\ell+1)$

$$\mathbb{P}\big(\hat{p}_{(m-i_0)} \leq x\big) = \mathbb{P}\Big(\sum_{i=1}^{m} \mathbb{1}\big[\hat{p}_i > x\big] \leq i_0\Big),$$

and then invoking the same arguments as above. □

Remark A1. A possible generalization would be to consider a rejection region on a vector of the ordered p-values, for instance considering two of the ordered p-values rather than just one. In such a case, deriving the upper bound would require evaluating the pairwise distribution of the ordered conformal p-values, and the p-value would depend on two hyperparameters rather than one.

Bates et al. (2023) showed a general result regarding the asymptotic distribution of test statistics of the form $\sum_{i=1}^{m} G(p_i)$, for some general class of functions G. In their work, this was used to formulate a correction to the Fisher combination test yielding a valid testing procedure using conformal p-values for the special case of $\pi_{\text{th}} = 0$. In the following theorem, we generalize their result to the setting of this paper, thereby paving the way for constructing more general test statistics which are valid asymptotically.

Theorem A2. *Let $T_G = \sum_{i=1}^{m} G(\hat{p}_i)$ be a test statistic for conformal p-values $\hat{p}_1, \ldots, \hat{p}_m$ and an increasing function $G : [0,1] \to [0,\infty)$ satisfying*

(i) $\int_0^1 G^{2+\eta}(u)du < \infty$;

(ii) $|\frac{1}{n-\ell+1}\sum_{j=1}^{n-\ell+1} G^k(j/(n-\ell+1)) - \int_0^1 G^k(u)du| = o(1/\sqrt{n-\ell})$, *for* $k \in \{1,2\}$;

(iii) $\max_{j\in\{1,\ldots,n-\ell+1\}} G(j/(n-\ell+1)) = o(\sqrt{n-\ell})$.

Then, under the null hypothesis $H_0 : \pi \leq \pi_{\text{th}}$, $\pi_{\text{th}} \in [0,1)$, if $m = \lfloor \gamma(n-\ell) \rfloor$ for some $\gamma > 0$, as $n-\ell \to \infty$

$$\mathbb{P}_{H_0}(T_G \leq r) \leq \pi_{\text{th}}^m + \sum_{k=1}^{m} \mathrm{B}_{\pi_{\text{th}}}^m(k) F_{G^k}\left(\frac{r + k(\sqrt{1+\gamma_k}-1)\int_0^1 G(u)\mathrm{d}u}{\sqrt{1+\gamma_k}}\right), \tag{A.2}$$

where F_{G^k} is the CDF of $\sum_{i=1}^{k} G(U_i)$, $U_i \overset{iid.}{\sim} \text{Unif}([0,1])$, and $\gamma_k = k/(n-\ell)$. Moreover, if $\mathbb{P}(G(\hat{p}_i) = 0, \forall i \in \bar{\mathcal{H}}_1) = 1$ and $\pi = \pi_{\text{th}}$ the inequality is exact.

Proof. By the law of total probability

$$\begin{aligned}\mathbb{P}_{H_0}\Big(\sum_{i=1}^{m} G(\hat{p}_i) \leq r\Big) &= \sum_{k=0}^{m} \mathbb{P}_{H_0}(k)\mathbb{P}\Big(\sum_{i=1}^{k} G(\hat{p}_i) + \sum_{i=k+1}^{m} G(\hat{p}_i) \leq r \mid \bar{\mathcal{H}}_0 = [k]\Big) \\ &\leq \sup_{\pi_0 \in [0,\pi_{\text{th}}]} \sum_{k=0}^{m} \mathrm{B}_{\pi_0}^m(k)\mathbb{P}\Big(\sum_{i=1}^{k} G(\hat{p}_i) \leq r \mid \bar{\mathcal{H}}_0 = [k]\Big),\end{aligned} \tag{A.3}$$

Now, by Theorem S1 of Bates et al. (2023), when $k = \lfloor \gamma_k(n-\ell) \rfloor$ we have that as $n-\ell \to \infty$

$$\mathbb{P}\Big(\sum_{i=1}^{k} G(\hat{p}_i) \leq r \mid \bar{\mathcal{H}}_0 = [k]\Big) \to q_k^*,$$

where q_k^* is such that $r = \sqrt{1+\gamma_k} Q_{G^k}(q_k^*) - k(\sqrt{1+\gamma_k}-1)\int_0^1 G(u)\mathrm{d}u$, and $Q_{G^k}(q_k^*)$ is the q_k^* quantile of $\sum_{i=1}^{k} G(U_i)$. Isolating q_k^* yields

$$q_k^* = F_{G^k}\left(\frac{r + k(\sqrt{1+\gamma_k}-1)\int_0^1 G(u)\mathrm{d}u}{\sqrt{1+\gamma_k}}\right).$$

It follows that

$$\lim_{n-\ell\to\infty} \mathbb{P}_{H_0}\Big(\sum_{i=1}^m G(\hat{p}_i) \le r\Big) \le \pi_{\text{th}}^m + \sum_{k=1}^m \mathrm{B}_{\pi_{\text{th}}}^m(k) F_{G^k}\left(\frac{r + k(\sqrt{1+\gamma_k}-1)\int_0^1 G(u)\mathrm{d}u}{\sqrt{1+\gamma_k}}\right),$$

where we use that q_k^* is a decreasing sequence. □

It follows from the theorem that an asymptotically valid p-value is

$$\hat{u}_G = \pi_{\text{th}}^m + \sum_{k=1}^m \mathrm{B}_{\pi_{\text{th}}}^m(k) F_{G^k}\left(\frac{T_G + k(\sqrt{1+\gamma_k}-1)\int_0^1 G(u)\mathrm{d}u}{\sqrt{1+\gamma_k}}\right). \tag{A.4}$$

Remark A3. We have altered the formulation slightly compared to Bates et al. (2023) since we want the function G to map into non-negative numbers and being an increasing function such that a small value of the test statistic is evidence against the null, thereby allowing the inequality of Eq. (A.3). Moreover, we want G to be such that $\mathbb{P}(G(\hat{p}_i) \le x, \forall i \in \bar{\mathcal{H}}_1)$ is large for small x, which is an important property to have a tight approximation in Eq. (A.3).

The following remarks, in combination with Theorem A2, provides the proof of Theorem 3.3.

Remark A4. A variant of the Fisher combination test uses $G(\hat{p}) = -2\log\big(\frac{1}{n-\ell+1}\big) + 2\log(\hat{p})$, thereby maintaining the shape of the typical Fisher combination function, while making it an increasing function. The result of Bates et al. (2023) can be exploited in this case by noticing that

$$\mathbb{P}\Big(-2k\log\Big(\frac{1}{n-\ell+1}\Big) + 2\sum_{i=1}^k \log(\hat{p}_i) \le r\Big) = 1 - \mathbb{P}\Big(-2\sum_{i=1}^k \log(\hat{p}_i) \le -2k\log\Big(\frac{1}{n-\ell+1}\Big) - r\Big).$$

Particularly, it holds that using this test statistic is equivalent to using the typical Fisher test statistic when $\pi_{\text{th}} = 0$. The technical conditions (i)-(iii) regarding the function G can be verified, see Remark S1 of Bates et al. (2023) for details.

Remark A5. A Sum combination test uses $G(\hat{p}) = \hat{p}$ for which $\int_0^1 G(u)\mathrm{d}u = 1/2$ and $G(U) \overset{d}{=} \text{Unif}([0,1])$, resulting in $\sum_{i=1}^k G(U_i) \overset{d}{=} \text{IH}(k)$ where $U_i \overset{iid.}{\sim} \text{Unif}([0,1])$ and $\text{IH}(k)$ is the Irwin-Hall distribution with parameter k. The technical conditions (i)-(iii) regarding the function G can be verified: condition (i) trivially holds since a polynomial of finite order is finite; for condition (ii) we note that $G(u)$ is increasing and $G'(u) = 1$, thus for $k \in \{1, 2\}$

$$\begin{aligned}
\Big|\frac{1}{n-\ell+1}\sum_{j=1}^{n-\ell+1}\frac{j^k}{(n-\ell+1)^k} - \int_0^1 u^k\mathrm{d}u\Big| &\le \sum_{j=1}^{n-\ell+1}\Big|\frac{1}{n-\ell+1}\frac{j^k}{(n-\ell+1)^k} - \int_{(j-1)/(n-\ell+1)}^{j/(n-\ell+1)} u^k\mathrm{d}u\Big| \\
&\le \sum_{j=1}^{n-\ell+1}\int_{(j-1)/(n-\ell+1)}^{j/(n-\ell+1)}\Big|\frac{j^k}{(n-\ell+1)^{k+1}} - u^k\Big|\mathrm{d}u \\
&\le \sum_{j=1}^{n-\ell+1}\frac{kj^{k-1}}{(n-\ell+1)^{k+1}} = \mathrm{O}\Big(\frac{1}{n-\ell}\Big)
\end{aligned}$$

where the first two inequalities are due to the triangle inequality, and the third inequality follows by the mean value theorem and the chain rule; finally condition (iii) immediately holds as $\max_{j\in\{1,\dots,n-\ell+1\}} G(j/(n-\ell+1)) = 1$.

A.1 Multiple testing with combination test p-values

For completeness of presentation we begin by defining the PRDS property (Benjamini & Yekutieli, 2001).

Definition A6. A set $\mathcal{D}$ is called non-decreasing if $x \in \mathcal{D}$ and $x \le y$ implies $y \in \mathcal{D}$.

Definition A7. For any non-decreasing set $\mathcal{D}$, and for each $i \in I_0$ such that $\mathbb{P}(X \in \mathcal{D}|X_i = x)$ is non-decreasing in x, then X is PRDS on I_0.

Theorem 3.5 states that the Storey conformal data contamination p-values are PRDS. This is a result motivated by Theorem 2.4 of Bates et al. (2023) stating that conformal p-values are PRDS. The implication of Theorem 3.5 is that we can use the BH procedure with the conformal data contamination p-values and maintain FDR control.

Proof of Theorem 3.5. Let F be a shorthand for the function in Eq. (2), i.e., $\hat{u}_k^{\text{storey}} = F(\hat{p}_1^k, \ldots, \hat{p}_{m_k}^k)$. It follows that F is entry-wise monotone increasing (not strictly), i.e., $F\big(\hat{p}_1^k, \ldots, (\hat{p}_j^k)', \ldots, \hat{p}_{m_k}^k\big) \geq F\big(\hat{p}_1^k, \ldots, \hat{p}_j^k, \ldots, \hat{p}_{m_k}^k\big)$ where $(\hat{p}_j^k)' \geq \hat{p}_j^k$ for any $j \in \{1, \ldots, m_k\}$. Let $Y = \big(\hat{u}_1^{\text{storey}}, \ldots, \hat{u}_K^{\text{storey}}\big)$ be the Storey conformal data contamination p-values on the test sets. Denote by $X_k = \big(\hat{p}_1^k, \ldots, \hat{p}_{m_k}^k\big)$ the conformal p-values on the k-th test set, and by $X = (X_1, \ldots, X_k)$ the total conformal p-values on the test sets.

Let $y \geq y'$ and let A be an increasing set. Conditioning on the k-th conformal data contamination p-value and using the law of total probability

$$\mathbb{P}(Y \in A|Y_k = y) = \sum_{x \in \mathcal{X}} \frac{\mathbb{P}(Y \in A, Y_k = y, X_k = x)}{\mathbb{P}(Y_k = y)},$$

where $\mathcal{X} = \{1/(n - \ell + 1), \ldots, 1\}^m$. Now, by conditional independence of $Y_{-k} = (Y_1, \ldots, Y_{k-1}, Y_{k+1}, \ldots, Y_K)$ and Y_k

$$\mathbb{P}(Y \in A|Y_k = y) = \sum_{x \in \mathcal{X}} \mathbb{P}(Y \in A|X_k = x)\mathbb{P}(X_k = x|Y_k = y).$$

We know that $\mathbb{P}(Y_k|X_k = x) = \mathbb{1}[Y_k = F(x)]$, and so $\mathbb{P}(X_k|Y_k = y)$ is only non-zero when $y = F(x)$. Define $S(y) = \{x \in \mathcal{X} : y = F(x)\}$.

$$\mathbb{P}(Y \in A|Y_k = y) = \frac{\sum_{x \in S(y)} \mathbb{P}(X \in B|X_k = x)}{|S(y)|},$$

where $B = \{\bar{x} \in \mathcal{X}^K : [F(\bar{x}_1), \ldots, F(\bar{x}_K)] \in A\}$. Since A is an increasing set, and F is an increasing function, B is also an increasing set. By Theorem 2.4 of Bates et al. (2023), conformal p-values are PRDS, and as a corollary to this, we have that $\mathbb{P}(X \in B|X_k = x) \geq \mathbb{P}(X \in B|X_k = x')$ where $x \succeq x'$ ($\succeq$ denotes entry-wise inequality).

We define some notations: let $S^{\text{vec}}(y)$ denote a vectorization of the set $S(y)$ such that $S^{\text{vec}}(y) = (\tilde{x}_1^{S(y)}, \ldots, \tilde{x}_{|S(y)|}^{S(y)})$ with $\tilde{x}_i^{S(y)} \in S(y)$, $\tilde{x}_i^{S(y)} \neq \tilde{x}_j^{S(y)}$, $i \neq j$, $i, j \in [|S(y))|]$; let also $\bar{S}^{\text{vec}}(y)$ denote the vectorization of $\bar{S}(y) \subset S(y)$; finally let $S^{\text{c}}(y)$ denote the complement of $S(y)$.

Assume initially that $|S(y')| = |S(y)|$. Then, $S(y)$ dominates $S(y')$ since F is an entry-wise increasing function, i.e., there exists a permutation σ on $[|S(y))|]$ such that $(\tilde{x}_{\sigma(1)}^{S(y)}, \ldots, \tilde{x}_{\sigma(|S(y)|)}^{S(y)}) \succeq (\tilde{x}_1^{S(y')}, \ldots, \tilde{x}_{|S(y)|}^{S(y')})$, and hence,

$$\mathbb{P}(Y \in A|Y_k = y) \geq \frac{\sum_{x' \in S(y')} \mathbb{P}(X \in B|X_k = x')}{|S(y')|} = \mathbb{P}(Y \in A|Y_k = y').$$

Now for the case $|S(y)| > |S(y')|$. Let $\bar{S}(y)$ be the subset of $S(y)$ with cardinality $|S(y')|$ which dominates $S(y')$ while minimizing $\sum_{x \in \bar{S}(y)} \mathbb{P}(X \in B|X_k = x)$, i.e., there exists a permutation σ on $[|S(y')|]$ such that

$(\tilde{x}^{\bar{S}(y)}_{\sigma(1)}, \ldots, \tilde{x}^{\bar{S}(y)}_{\sigma(|S(y')|)}) \succeq (\tilde{x}^{S(y')}_1, \ldots, \tilde{x}^{S(y')}_{|S(y')|})$, and hence,

$$\begin{aligned}
\mathbb{P}(Y \in A | Y_k = y) &= \frac{\sum_{x \in \bar{S}(y)} \mathbb{P}(X \in B | X_k = x) + \sum_{x \in \bar{S}^c(y)} \mathbb{P}(X \in B | X_k = x)}{|S(y')| + |\bar{S}^c(y)|} \\
&\geq \frac{\sum_{x \in \bar{S}(y)} \mathbb{P}(X \in B | X_k = x)}{|S(y')|} \\
&\geq \frac{\sum_{x' \in S(y')} \mathbb{P}(X \in B | X_k = x)}{|S(y')|} \\
&= \mathbb{P}(Y \in A | Y_k = y').
\end{aligned}$$

Finally for the case $|S(y')| > |S(y)|$. Let $\bar{S}(y')$ be the subset of $S(y')$ with cardinality $S(y)$ which is dominated by $S(y)$ while maximizing $\sum_{x' \in \bar{S}(y')} \mathbb{P}(X \in B | X_k = x')$, i.e., there exists a permutation σ on $[|S(y)|]$ such that $(\tilde{x}^{\bar{S}(y')}_{\sigma(1)}, \ldots, \tilde{x}^{\bar{S}(y')}_{\sigma(|S(y)|)}) \preceq (\tilde{x}^{S(y)}_1, \ldots, \tilde{x}^{S(y)}_{|S(y)|})$. Now observe that

$$\begin{aligned}
\mathbb{P}(Y \in A | Y_k = y) &= \frac{\sum_{x \in S(y)} \mathbb{P}(X \in B | X_k = x)}{|S(y)|} \\
&\geq \frac{\sum_{x' \in \bar{S}(y')} \mathbb{P}(X \in B | X_k = x')}{|S(y)|} \\
&\geq \frac{\sum_{x' \in \bar{S}(y')} \mathbb{P}(X \in B | X_k = x') + \sum_{x' \in \bar{S}^c(y')} \mathbb{P}(X \in B | X_k = x')}{|S(y')|} \\
&= \mathbb{P}(Y \in A | Y_k = y').
\end{aligned}$$

The existence of such $\bar{S}(y)$ and $\bar{S}(y')$ is confirmed for the Storey conformal data contamination p-value. Consider initially that the Storey test statistic, denoted by T here, is $T = t + 1$ and $T' = t$, i.e., just a difference of one. Now, $S(y) = \{x \in \mathcal{X} : x_{(1)} \leq \cdots \leq x_{(K-t-1)} \leq \lambda < x_{(K-t)} \leq \cdots \leq x_{(K)}\}$, and $S(y') = \{x' \in \mathcal{X} : x'_{(1)} \leq \cdots \leq x'_{(K-t)} \leq \lambda < x'_{(K-t+1)} \leq \cdots \leq x'_{(K)}\}$, and so the only difference occurs at $x_{(K-t)}$ which for y is greater than or equal to λ while for y' it is less than or equal to λ. Now, for the case of $|S(y)| > |S(y')|$, for any $x' \in S(y')$ we can find an $x \in S(y)$ such that $x \succeq x'$. Taking for each $x' \in S(y')$ a unique $x \in S(y)$ with smallest $\mathbb{P}(X \in B | X_k = x)$ (avoiding using the same x more than once), we have constructed $\bar{S}(y)$. In the same way we can construct $\bar{S}(y')$, taking for each $x \in S(y)$ a unique $x' \in S(y')$ with the largest $\mathbb{P}(X \in B | X_k = x')$. Extending the argument to differences in the Storey test statistic of more than one is immediate. □

We have only proven the PRDS property for the Storey conformal data contamination p-values in Theorem 3.5, however, we expect the result can also be shown for the other conformal data contamination p-values.

B Additional discussion of the proposed data sharing procedure

We outline in Tables A1 and A2 all scenario variables and hyperparameters. Here we recall the notation for each variable, and discuss how it influences the collaborative data sharing method proposed in this work.

C Additional numerical experiments: Gaussian data

In this section, we present a simulation study with null distribution $P_0 \equiv \mathcal{N}(\mathbf{0}_2, \boldsymbol{I}_{2\times 2})$ and alternative $P_1 \equiv \mathcal{N}(\mu_1 \mathbf{1}_2, \boldsymbol{I}_{2\times 2})$. As conformal score we use $\hat{s}(X) = -\|X\|$ as it is a natural choice in light of the distribution P_0 (no learning of the score function is used in this case and so $\ell = 0$ throughout this section). We simulate $n = 200$ data points from P_0 to be the calibration data set and let the test data consist of $m = 50$ data points. The test data on average consists of $m(1 - \pi)$ data points sampled from P_0 while the remaining data points are sampled from P_1.

Table A1: Scenario variables.

Parameter	Description & Influence
n	The size of the null sample. As $n-\ell$ increases, the empirical distribution better approximates the true distribution, yielding more accurate conformal p-values. Moreover, this variable controls the smallest possible p-value, thereby controlling how much evidence one outlier can yield towards rejection in the conformal data contamination tests.
m	The number of test data points. As this variable increases, more evidence against the null can be aggregated yielding a more powerful conformal data contamination test.
π	The true contamination factor. As the difference $\pi - \pi_{\text{th}}$ increases, the power of the conformal data contamination tests increases.
K	The number of other data agents.
K_0	The number of other data agents satisfying the null hypothesis $H_0^k : \pi_k \leq \pi_{\text{th}}$. As this grows, the power of the multiple testing procedure decreases.

Table A2: Hyperparameters.

ℓ	The number of null data points used for learning the conformal score. As this variable increases the conformal score will tend to better separate outliers from inliers thereby improving outlier detection and increasing power of the conformal data contamination tests.
$\hat{s}$	The conformal score function. Choosing a good conformal score is key to achieving good separation between outliers and inliers, which also directly impacts the power and tightness of the conformal data contamination tests.
π_{th}	The chosen threshold on the contamination factor defining the null hypotheses. This hyperparameter controls the amount of contamination that is tolerated.
α	The significance level for the conformal data contamination tests. This hyperparameter controls the amount of evidence needed before concluding that the data from another agent is more contaminated than the threshold π_{th} allows.
γ	The hyperparameter in Storey's BH procedure. Controls the bias-variance trade-off for estimating K_0.
β	The significance level in the conformal outlier detection. This hyperparameter controls how much evidence is needed to conclude that a data point is an outlier. As this increases, the number of data points which is subsequently used for model fitting increases, with the risk of admitting more outliers.
ζ	The hyperparameter in Storey's BH procedure for conformal outlier detection. Controls the bias-variance trade-off for estimating the number of inliers among the acquired data.
λ	The hyperparameter in the Storey test statistic.
i_0	The hyperparameter in the Quantile test statistic.

C.1 Hyperparameters of the conformal data contamination tests

We compare, in Figure A1, the power estimated using 10000 simulations when $\alpha = 0.05$, $\mu_1 = 4$, $\pi = 0.7$, $\pi_{\text{th}} = 0.5$, $n = 200$, and $m = 50$, depending on the hyperparameters of the conformal data contamination tests, for the two cases $\mu_1 = 2$ and $\mu_1 = 4$. We can observe that for appropriate hyperparameter choices, the Storey and Quantile conformal data contamination tests can achieve the highest power. Moreover, when the outliers are easily separated from the inliers ($\mu_1 = 4$), λ should be chosen as a very small value and the choice of i_0 should also be sufficiently small, meanwhile for $\mu_1 = 2$, λ and i_0 should be chosen moderately. This highlights a weakness of the Storey and Quantile tests as they are sensitive to the hyperparameter choice. Note also that the Fisher test is relatively powerful when the outliers are easily separated, however it is very weak when this is not the case. This is a sensible observation when considering the function used in the Fisher test statistic as it emphasizes the information in very small p-values, but when the outliers are not easily separated, the p-values for the outliers will not be very small and so is less emphasized.

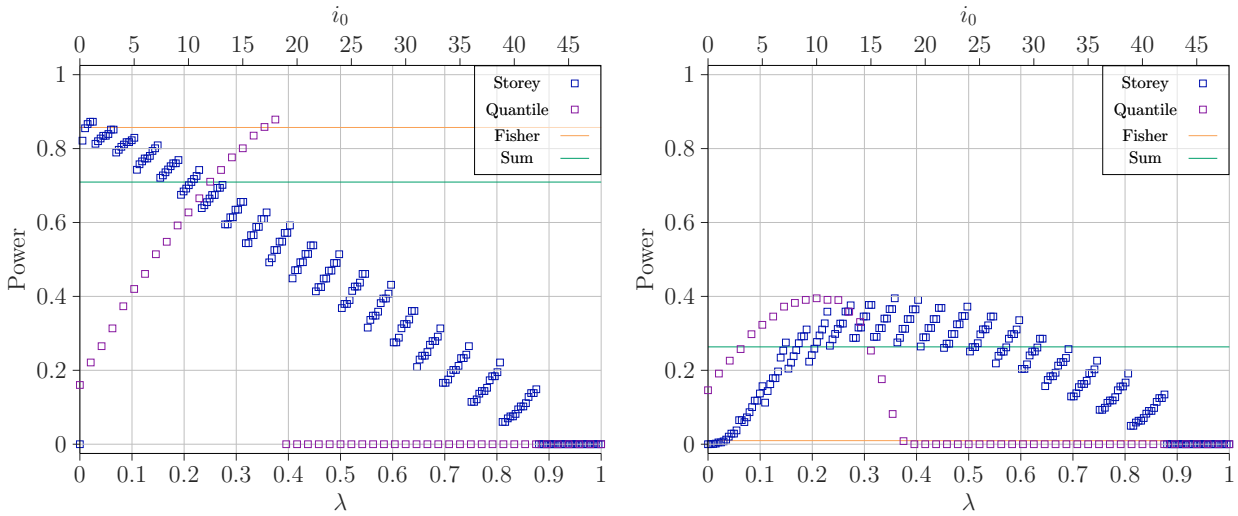

(a) $\mu_1 = 4$.

(b) $\mu_1 = 2$.

Figure A1: Comparison of power of the different tests for varying hyperparameters.

C.2 Special case of classical two-sample testing

In the particular case of $\pi_{\text{th}} = 0$, the problem reverts to the classical two-sample testing problem. This can be solved using classical techniques such as Kolmogorov-Smirnov and Cramér-von Mises, but also permutation tests as in Konstantinou et al. (2024), as well as the combination tests of Bates et al. (2023). Notably, the Fisher conformal data contamination test proposed in this work coincides with the Fisher combination test of Bates et al. (2023) for $\pi_{\text{th}} = 0$.

The setup is the following: we observe two samples of data $\mathcal{D}_0 = \{X_i\}_{i=1}^n$ and $\mathcal{D}_1 = \{X_i\}_{i=n+1}^m$. Now, many different two-sample tests can be constructed by defining a summary statistic $\hat{s}(X)$, and then comparing the empirical distributions of the summary statistic on the two samples: $\hat{F}_0 = \hat{F}(\hat{s}(\mathcal{D}_0))$ and $\hat{F}_1 = \hat{F}(\hat{s}(\mathcal{D}_1))$, where $\hat{s}(\mathcal{D}_1) = \{\hat{s}(X_{n+i})\}_{i=1}^m$ and $\hat{F}$ computes the empirical CDF. Now, a test statistic can be constructed from these two empirical distributions $T(\hat{F}_0, \hat{F}_1)$.

Kolmogorov-Smirnov test: A classic test statistic is the Kolmogorov-Smirnov statistic

$$T_{\text{KS}}(\hat{F}_0, \hat{F}_1) = \sup_x |\hat{F}_0(x) - \hat{F}_1(x)|,$$

for which asymptotically as $n, m \to \infty$

$$\mathbb{P}\left(T_{\text{KS}}(\hat{F}_0, \hat{F}_1) > \sqrt{-\ln(\alpha/2)\frac{n+m}{2nm}}\right) \leq \alpha,$$

with significance level $\alpha \in (0, 1)$.

Permutation tests: Permutation tests are based on constructing random permutations by randomly assigning data points from $\mathcal{D}_0$ and $\mathcal{D}_1$ to new splits $\mathcal{D}_{i,0}$ and $\mathcal{D}_{i,1}$ for $i = 1, \ldots I$ where I is the number of permutations. The idea is that under the null, the test statistic $T(\hat{F}_0, \hat{F}_1)$ is exchangeable with $T_i = T(\hat{F}(\hat{s}(\mathcal{D}_{i,0})), \hat{F}(\hat{s}(\mathcal{D}_{i,1})))$ for $i = 1, \ldots I$ (Konstantinou et al., 2024). This allows for computing a p-value based on the empirical distribution of the test statistic among the permutations, allowing for flexibility in the choice of test statistic T, however, at a computational cost. We consider as examples test statistic using the Kolmogorov-Smirnov difference, the L_2-norm, and a test statistic using one class support vector machines (OCSVM). For the OCSVM, a number of training permutations are made yielding $T_1, \ldots, T_B$, followed by a number of calibration permutations $T_{B+1}, \ldots, T_I$. Then, the OCSVM is fitted to $T_1, \ldots, T_B$, and the output scores of the OCSVM on $T_{B+1}, \ldots, T_I$ are used a calibration data with the OCSVM score of $T(\hat{F}_0, \hat{F}_1)$ being the observed test statistic.

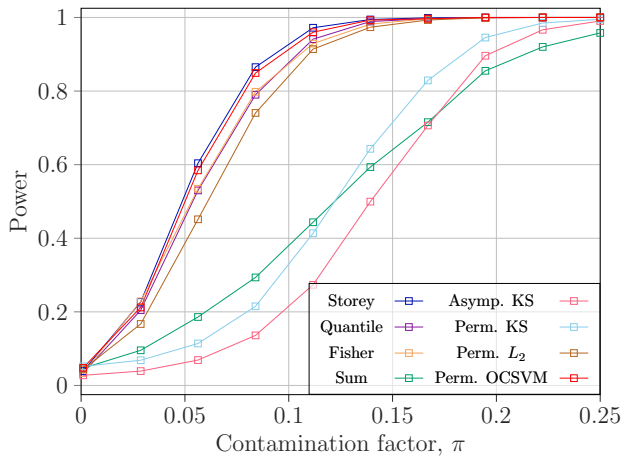

Figure A2: Comparison of the power of different two-sample tests for varying contamination factors.

Numerical power comparison: A power comparison across 2000 simulations of different two-sample tests for varying contamination factors when $\alpha = 0.05$, $n = 200$, $m = 100$, $\mu_1 = 4$, $\lambda = \lfloor n/32 \rfloor/(n+1)$, and $i_0 = 5$, is shown in Figure A2. We have used 200 permutations for fitting OCSVM (in Perm. OCSVM), and 500 permutations for calibration (in the permutation tests). Figure A2 indicates that the weakest tests are the Kolmogorov-Smirnov tests and the Sum conformal data contamination test, and the other five tests achieve relatively similar power. We note that it intuitively makes sense that the Fisher combination test is powerful here as this test statistic emphasizes the contribution of very small p-values more so than the other test statistics, which is a desirable property when looking for a minimum of just one outlier. On the other hand, the Sum conformal data contamination test is relatively weak in this setting. These properties of these two tests were also discussed in the context of Figure A1. Our main take-away is that the proposed conformal data contamination tests are competitive with state-of-the-art tests.

C.3 Ablation study: test and calibration data size

In this section we study the performance of the conformal data contamination tests for varying parameter settings. Specifically, the focus will be on varying contamination factors, π, calibration data sizes, $n - \ell$, and test data sizes, m. We set $\lambda = \lfloor (n-\ell)/12 \rfloor/(n-\ell+1)$, $i_0 = \lfloor m/1.5 \rfloor$, $\alpha = 0.05$, $\pi_{\text{th}} = 0.1$, and $\mu_1 = 4$.

Consider an ablation study for the contamination factor, π, and the test data size, m, with $n = 100$. We show in Figure A3a the power estimated using 1000 simulations for the conformal data contamination tests. We observe that the power relatively quickly goes to 1 as the contamination factor increases, and there is a general tendency that the power increases with m, which is also to be expected. In this simulation study, the Storey conformal data contamination test achieves the highest power.

The influence of the calibration data size, $n - \ell$, with $m = 100$ and $\pi = 0.3$, is visualized in Figure A3b. Here we show the power estimated using 10000 simulations for the conformal data contamination tests. The Storey and Quantile conformal data contamination tests have the highest power, achieving a power near 1 for $n - \ell \geq 20$. This highlights that powerful conformal data contamination tests are possible even for extremely modest calibration data sizes.

With the insights gained through these numerical experiments we can highlight an advantage of the conformal data contamination tests. Conformal outlier detection relies on having a calibration data size, $n - \ell$, relatively large compared to the test data size, m, see Mary & Roquain (2022). Meanwhile, conformal data contamination tests only improve with increasing m, and so are more convenient in settings with small $n - \ell$ and large m.

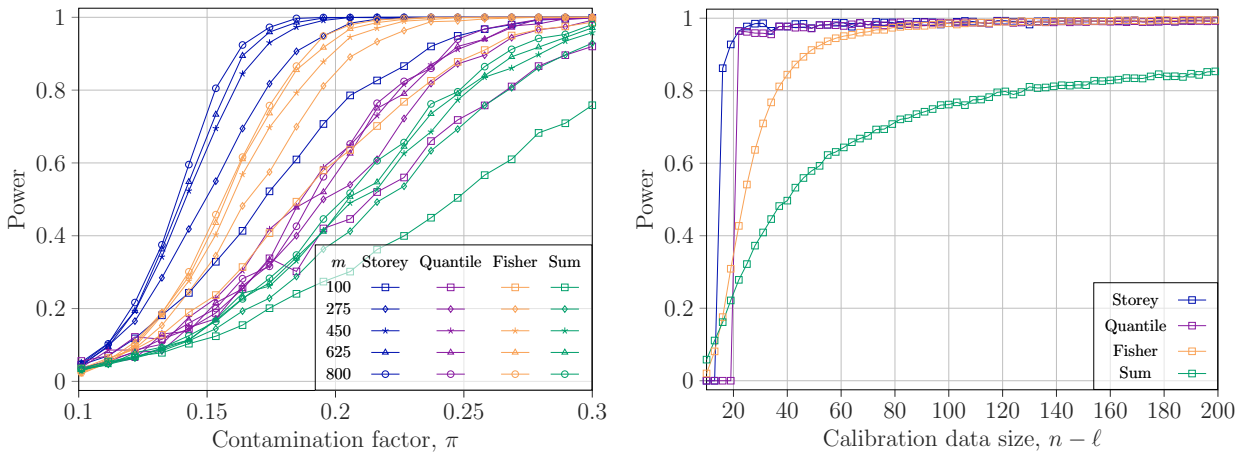

(a) (b)

Figure A3: Ablation study for varying contamination factors, π, calibration data sizes, $n - \ell$, and test data sizes, m.

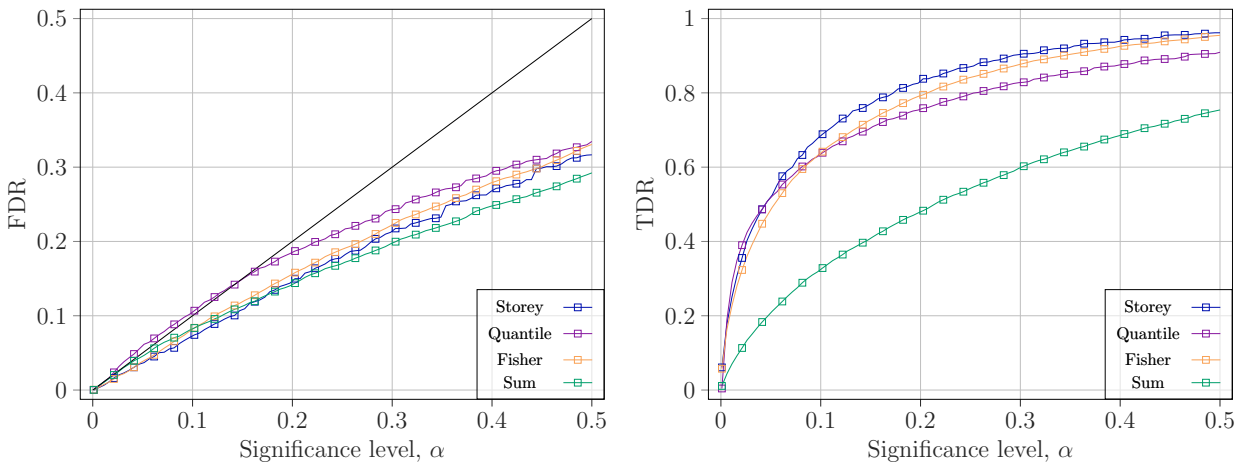

(a) FDR. (b) TDR.

Figure A4: Comparison of FDR and TDR estimated of the different tests.

C.4 Multiple testing with conformal data contamination p-values

Consider the same scenario as before except now $K = 20$ agents are sending test data for which $K_0 = 10$ of them have $\pi_0 = \pi_{\text{th}}$ and the remaining $K - K_0 = 10$ have $\pi_1 > \pi_{\text{th}}$. We use Storey's BH procedure on the sequence of conformal data contamination p-values with Storey's hyperparameter $\gamma = 0.5$.

In Figure A4 we show FDR and TDR curves estimated by 2000 simulations for the four different conformal data contamination tests presented in Section 3, when $\alpha = 0.05$, $\mu_1 = 4$, $n = 200$, $m = 100$, $\pi_{\text{th}} = 0.2$, $\pi_1 = 0.3$, $\lambda = \lfloor (n - \ell)/32 \rfloor / (n - \ell + 1)$, and $i_0 = \lfloor m/1.5 \rfloor$. We notice from Figure A4a that all the methods are conservative. In terms of power Figure A4b shows that the Storey conformal data contamination test achieves the highest power, with comparable power between Quantile and Fisher test, and the Sum test tends to have the lowest power. We emphasize that this does not mean that in general the Sum test is the worst alternative among those considered here. Which of the proposed tests performs the best depends on many factors, herein π_{th} and the conformal score.

D Additional numerical experiments: Real datasets

In this section, we provide extensive numerical results on a variety of data contamination scenarios with real image classification datasets. We consider several conformal scores and evaluate both the outlier detection capability, the performance of the proposed conformal data contamination tests, and the performance on the relevant classification task for the proposed data acquisition procedure as well as relevant baselines.

In the following we list the data contamination scenarios.

(S1) Mixture of retinal fundus images (color) from RetinaMNIST (Liu et al., 2022) and EyePACS (Dugas et al., 2015; Gulshan et al., 2016), as also considered in Section 5. The classification task is to distinguish between five degrees of severity of diabetic retinopathy ranging from *no diabetic retinopathy* to *proliferate retinopathy*.

(S2) Mixture of optical coherence tomography (OCT) images (grayscale) taken from University of California San Diego found in OCTMNIST (Kermany et al., 2018) and OCT images from Noor Eye Hospital (NEH) (Sotoudeh-Paima et al., 2022). The classification task is to distinguish *choroidal neovascularization* from *normal retina*.

(S3) Mixture of microscopic blood smear images (color) from BloodMNIST (Acevedo et al., 2020) and (Bodzas et al., 2023). The classification task to to distinguish between the labels *basophile*, *eosinophile*, *lymphocyte*, and *monocyte*.

(S4) Mixture of images (grayscale) of lower- and uppercase handwritten letters A, B, and C from FEMNIST (Caldas et al., 2019).

(S5) Mixture of noise-free and noisy images (grayscale) of handwritten digits 1, 4, and 7 from MNIST (Lecun et al., 1998). We sample the feature noise as a Gaussian vector with zero mean and covariance kernel $\varsigma^2 \exp(-\varrho \|p_i - p_j\|)$ for pixel positions $p_i = i$ and $p_j = j$ for the i-th row and j-column in the pixel image, and add it to the image. After adding noise, each image is normalized to the interval $[0, 1]$. We set $\varsigma = 0.3$ and $\varrho = 0.1$.

(S6) Mixture of images (grayscale) of handwritten digits 1, 4, and 7 from MNIST with and without noisy labels.

(S7) Mixture of images of dogs, deer, and cats from CIFAR-10 and CIFAR-10C contaminated with brightness (Hendrycks & Dietterich, 2019).

(S8) Mixture of images of dogs, deer, and cats from CIFAR-10 and CIFAR-10C contaminated with snow (Hendrycks & Dietterich, 2019).

(S9) Mixture of images of classes 'n01641577', 'n04562935', 'n02233338', 'n09246464', and 'n03992509' from Tiny-ImageNet and Tiny-Imagenet-C contaminated with defocus blur (Hendrycks & Dietterich, 2019).

The images in (S1), (S2), and (S3) are cropped and down sampled in the same way to 28×28 ($28 \times 28 \times 3$ for color images) pixels. All images are pixel-wise normalized to the interval $[0, 1]$.

In all cases we sample randomly n in-distribution datapoints for the 0-th data agent. Meanwhile, the k-th data agent observes $2m$ datapoints and the number of OOD datapoints is sampled from $\mathrm{B}^{2m}_{\pi_k}$. We sample the contamination factors iid according to a standard uniform distribution, such that approximately half of the test data are outliers. For all the numerical experiments we set $n = 100$, $\ell = 60$, $m = m_{r0} = m_{r1} = 40$, $K = 10$, $\lambda = \lfloor (n - \ell)/8 \rfloor / (n - \ell + 1)$, $i_0 = \lfloor m/3 \rfloor$, and $\gamma = 0.5$. Here we chose n relatively small as this is the regime where data sharing is most interesting, and we decided on a 60-40 split between data for fitting conformal scores and calibration data, as we found that having enough data for fitting the conformal score is often critical. The hyperparameters λ and i_0 were chosen in light of the insights gained in the ablation study in Section C. The hyperparameter for Storey's BH procedure, γ, was chosen somewhat arbitrarily at the intermediate value 0.5.

We consider different conformal scores:

(C1) OCSVM (Schölkopf et al., 2001).

(C2) Class-wise score based on OCSVM.

(C3) Isolation forest (IF) (Liu et al., 2008).

(C4) Class-wise score based on IF .

(C5) Autoencoder (AE) (Hawkins et al., 2002).

(C6) Class-wise score based on AE.

(C7) PU classification using LR (Marandon et al., 2024).

(C8) Class-wise score based on PU classification using LR.

(C9) PU classification using SVC (Marandon et al., 2024).

(C10) Class-wise score based on PU classification using SVC.

(C11) PU classification using LR on principal components explaining at least 80 % of the variance.

(C12) Class-wise score based on PU classification using LR on principal components explaining at least 80 % of the variance.

(C13) PU classification using a Residual Network with four layers (ResNet-4) (He et al., 2016).

(C14) Class-wise score based on PU classification using a ResNet-4.

Class-wise scores are defined as

$$\hat{s}((X,Y),(Z_1,\ldots,Z_\ell),(Z_{\ell+1},\ldots,Z_n,Z_1^1,\ldots,Z_m^K)) = \sum_{i\in\mathcal{Y}} \mathbb{1}[Y=i]\ \hat{s}_i(X, X_{\text{train}}^{Y=i}, X_{\text{cal+test}}^{Y=i}),$$

where $\mathcal{Y}$ denotes the set of labels, $X_{\text{train}}^{Y=i} = (X_j : Y_j = i, j \in [\ell])$ denotes the subset of $X_1,\ldots,X_\ell$ for which $Y_j = i$ for $j \in [\ell]$, and likewise $X_{\text{cal+test}}^{Y=i} = (X_{\ell+j} : Y_{\ell+j} = i, j \in [n-\ell]) \cup (X_j^k : Y_j^k = i, j \in [m], k \in [K])$.

Scores (C1)-(C4) and (C7)-(C12) are implemented using the standard implementations in `scikit-learn`, the scores (C5) and (C6) are implemented in `tensorflow`, and the scores (C13) and (C14) are implemented in `pytorch`. Details for the implementation of the score functions are given in Section E.

D.1 Marginal power and type I error probability study

To compare the different conformal scores on the various data contamination scenarios we will consider the marginal power of the conformal outlier detection and conformal data contamination tests, as well as the type I error probability for the conformal data contamination tests, through the empirical CDFs of the conformal p-values and the conformal data contamination p-values, respectively. Particularly, we will use the area under the curve (AUC) to summarize the CDFs to scalars.

For conformal outlier detection, $\text{AUC}^{\text{cod}} = 1/(n-\ell+2)\sum_{i=1}^{n-\ell+1} \hat{F}_{\text{outlier}}(i/(n-\ell+1))$ where $\hat{F}_{\text{outlier}}$ is the empirical CDF of the conformal p-value among outliers. AUC^{cod} ranges from 0 to 1 with 0.5 corresponding to random guessing and 1 corresponding to perfect detection.

For conformal data contamination tests, $\text{AUC}^{cdct}_{(\pi,\pi_{\text{th}})} = \int_0^1 \hat{F}_{(\pi,\pi_{\text{th}})}(x)\mathrm{d}x$ where $\hat{F}_{(\pi,\pi_{\text{th}})}$ is the empirical CDF of the conformal data contamination test p-value for data contamination threshold π_{th} and when the data contamination level is π. For comparing the power we are interested in $\pi > \pi_{\text{th}}$ and a larger $\text{AUC}^{cdct}_{(\pi,\pi_{\text{th}})}$ the better, while for numerically evaluating the conservativeness we are interested in $\pi = \pi_{\text{th}}$ which should be less than but also close to 0.5.

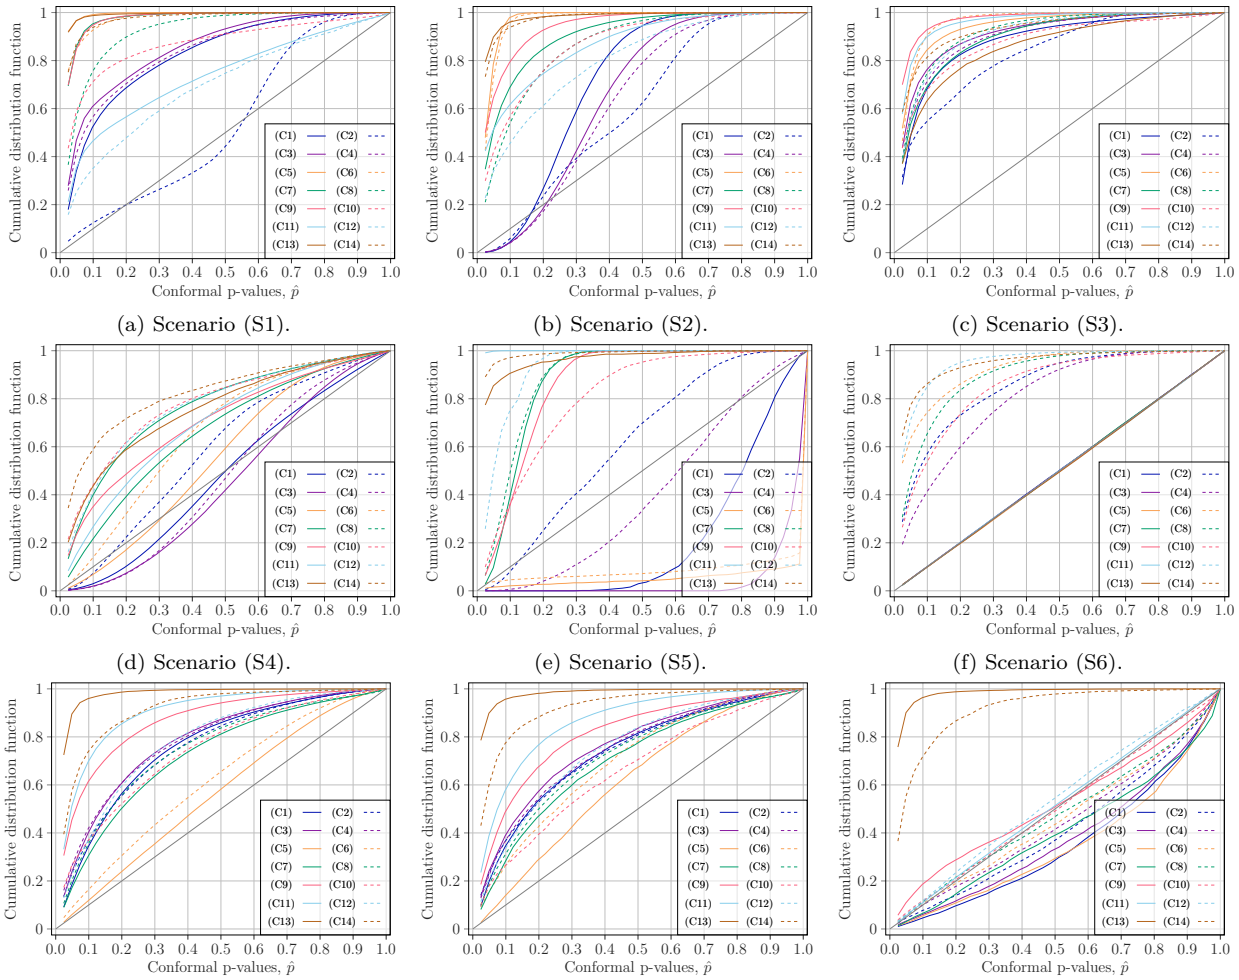

(a) Scenario (S1). (b) Scenario (S2). (c) Scenario (S3).

(d) Scenario (S4). (e) Scenario (S5). (f) Scenario (S6).

(g) Scenario (S7). (h) Scenario (S8). (i) Scenario (S9).

Figure A5: *Power study:* Empirical CDFs of AUC^{cod} for the different conformal scores and data contamination scenarios.

We report in Table A3 AUC^{cod} for the considered data contamination scenarios and for each of the different conformal scores, and show the full CDFs in Figure A5. Immediately, we can notice that the adaptive scores using the PU classification techniques of Marandon et al. (2024) consistently perform well. For all the medical data scenarios (S1), (S2), and (S3), AdaDetect using SVC or ResNet-4, i.e. (C9) or (C13), as well as the AE, i.e. (C5), perform very well. For scenarios (S4), (S5), and (S6), class-wise scores with AdaDetect generally perform well, i.e. (C8), (C10), (C12), and (C14), consistently yielding accurate outlier detection. Meanwhile, the most effective score for the CIFAR-10 scenarios (S7) and (S8) is AdaDetect using ResNet-4, i.e. (C13), with LR on the principal components, i.e. (C11), not far behind. Finally, the only conformal score among those considered that effectively detects outliers in the Tiny-ImageNet scenario (S9) is AdaDetect using ResNet-4, i.e. (C13) and (C14).

These results may suggest that AdaDetect using a Residual Network is a good choice for image data, but such a neural network approach is of higher complexity than for instance AdaDetect with LR and can also suffer from overfitting during the training phase if the training algorithm is not well specified, leading to catastrophic failure of outlier detection. For a given data application, it is crucial to choose a high performing conformal score which may require some domain knowledge.

While these results for varying data contamination scenarios with different conformal scores are not exhaustive, they give an indication of the achieveable outlier detection performance, and provide some guidelines for

Table A3: *Power study:* AUC for conformal outlier detection computed from 2000 simulations. The best conformal score for each data contamination scenario is highlighted with boldface.

	(S1)	(S2)	(S3)	(S4)	(S5)	(S6)	(S7)	(S8)	(S9)
(C1)	0.8164	0.7110	0.8694	0.4819	0.2246	0.5011	0.7531	0.7279	0.3536
(C2)	0.5343	0.6047	0.8240	0.5852	0.6158	0.8451	0.7504	0.7292	0.4128
(C3)	0.8439	0.6582	0.9003	0.4539	0.0586	0.4996	0.7738	0.7501	0.3774
(C4)	0.8261	0.6333	0.8894	0.4773	0.4077	0.7979	0.7734	0.7361	0.4426
(C5)	**0.9726**	0.9562	0.9287	0.5554	0.0712	0.4984	0.5594	0.6116	0.3514
(C6)	0.9581	0.9515	0.9450	0.6569	0.0997	0.9088	0.6092	0.6913	0.4542
(C7)	0.9584	0.8921	0.8863	0.6580	0.8676	0.4999	0.7164	0.6948	0.4046
(C8)	0.9105	0.8536	0.9058	0.7572	0.8792	0.8850	0.7409	0.7124	0.4704
(C9)	0.9580	0.9259	**0.9537**	0.6979	0.8525	0.4992	0.8543	0.7999	0.5160
(C10)	0.8571	0.8549	0.8645	0.7678	0.8019	0.8438	0.7345	0.6530	0.4999
(C11)	0.7283	0.8568	0.9435	0.6903	**0.9759**	0.4995	0.8896	0.8502	0.5110
(C12)	0.6904	0.7917	0.8982	0.7633	0.9223	**0.9357**	0.7725	0.7401	0.5326
(C13)	0.9709	**0.9613**	0.8478	0.7502	0.9479	0.4978	**0.9608**	**0.9610**	**0.9621**
(C14)	0.9565	0.9578	0.9156	**0.8145**	0.9666	0.9331	0.9003	0.9063	0.8976

Table A4: *Power study:* AUC for conformal data contamination tests computed from 2000 simulations when $\pi = 0.3$. The best test statistic for each data contamination scenario is highlighted with boldface.

		(S1)			(S2)			(S3)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	**0.9514**	**0.8233**	0.5867	**0.8607**	0.6470	0.3857	**0.9462**	**0.8108**	0.5720
	Quantile	0.8055	0.7309	**0.6279**	0.7969	**0.7185**	**0.6110**	0.7998	0.7244	**0.6207**
	Fisher	0.9289	0.7351	0.4456	0.8474	0.5762	0.2815	0.9245	0.7321	0.4446
	Sum	0.8987	0.7634	0.5632	0.8575	0.6952	0.4775	0.8959	0.7589	0.5558
		(S4)			(S5)			(S6)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	**0.7866**	0.5165	0.2594	0.7636	0.5989	0.4263	**0.9140**	**0.7367**	0.4780
	Quantile	0.7206	**0.6316**	**0.5175**	0.8170	**0.7451**	**0.6445**	0.7910	0.7140	**0.6091**
	Fisher	0.7514	0.4292	0.1704	0.8247	0.5416	0.2606	0.8954	0.6594	0.3585
	Sum	0.7668	0.5674	0.3485	**0.8786**	0.7295	0.5191	0.8759	0.7245	0.5123

choosing conformal scores. In the following, we limit our attention to class-wise AdaDetect with LR, i.e. (C8), as this score is simple and perform consistently well for scenarios (S1)-(S6), which we intermediately restrict our attention to.

In Table A4, we display $\text{AUC}^{\text{cdct}}_{(\pi,\pi_{\text{th}})}$, for $\pi > \pi_{\text{th}}$ thereby indicating the power of the tests, for the four different proposed conformal data contamination tests for varying values of the data contamination threshold, π_{th}, when the data contamination level is $\pi = 0.3$. The AUC decreases for increasing data contamination thresholds, which is to be expected as distinguishing between smaller differences in data contamination is naturally harder. We notice that in almost all cases, the Storey and Quantile tests, which are valid non-asymptotically, perform better than the Fisher and Sum tests, which are only valid asymptotically. Further, we observe that the Storey test usually has the highest AUC for small data contamination thresholds, whereas the Quantile test is comparatively better for larger data contamination thresholds. A similar observation holds when comparing the Fisher and Sum tests, with the Fisher test performing better at small data contamination thresholds, and vice versa, as to be expected.

In Table A5, we display $\text{AUC}^{\text{cdct}}_{(\pi,\pi_{\text{th}})}$, for $\pi = \pi_{\text{th}}$ thereby indicating the probability of false rejection, again for all the conformal data contamination tests for varying values of $\pi = \pi_{\text{th}}$. As the developed theory promises,

Table A5: *Type I error probability study:* AUC for conformal data contamination tests computed from 2000 simulations when $\pi = \pi_{\text{th}}$.

		(S1)			(S2)			(S3)		
		$\pi_{\text{th}}=0$ $\pi=0$	$\pi_{\text{th}}=0.1$ $\pi=0.1$	$\pi_{\text{th}}=0.2$ $\pi=0.2$	$\pi_{\text{th}}=0$ $\pi=0$	$\pi_{\text{th}}=0.1$ $\pi=0.1$	$\pi_{\text{th}}=0.2$ $\pi=0.2$	$\pi_{\text{th}}=0$ $\pi=0$	$\pi_{\text{th}}=0.1$ $\pi=0.1$	$\pi_{\text{th}}=0.2$ $\pi=0.2$
(C8)	Storey	0.4906	0.4165	0.3537	0.4990	0.3358	0.2402	0.4934	0.4069	0.3517
	Quantile	0.4796	0.4932	0.4908	0.5035	0.5033	0.4832	0.5150	0.4987	0.4933
	Fisher	0.3843	0.3112	0.2299	0.4110	0.2537	0.1565	0.4106	0.3033	0.2383
	Sum	0.4184	0.4082	0.3591	0.4519	0.3829	0.3168	0.4545	0.4021	0.3669
		(S4)			(S5)			(S6)		
		$\pi_{\text{th}}=0$ $\pi=0$	$\pi_{\text{th}}=0.1$ $\pi=0.1$	$\pi_{\text{th}}=0.2$ $\pi=0.2$	$\pi_{\text{th}}=0$ $\pi=0$	$\pi_{\text{th}}=0.1$ $\pi=0.1$	$\pi_{\text{th}}=0.2$ $\pi=0.2$	$\pi_{\text{th}}=0$ $\pi=0$	$\pi_{\text{th}}=0.1$ $\pi=0.1$	$\pi_{\text{th}}=0.2$ $\pi=0.2$
(C8)	Storey	0.5065	0.3045	0.1676	0.5011	0.3524	0.2746	0.5084	0.3888	0.3078
	Quantile	0.5068	0.4746	0.4374	0.5043	0.5022	0.4904	0.4996	0.5033	0.4930
	Fisher	0.4194	0.2168	0.1029	0.4078	0.2427	0.1412	0.4117	0.2900	0.2049
	Sum	0.4570	0.3441	0.2490	0.4492	0.3906	0.3333	0.4460	0.3997	0.3496

Table A6: TDR computed from 2000 simulations at significance level $\alpha = 0.1$ using Storey's BH procedure with parameter $\gamma = 0.5$. Data contamination factors are simulated as $\pi_k \sim \text{Uniform}([0, 1])$. The best test statistic for each data contamination scenario is highlighted with boldface.

		(S1)			(S2)			(S3)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	**0.9128**	**0.8147**	0.7033	0.8014	0.5995	0.4145	**0.9122**	**0.8056**	0.6838
	Quantile	0.8104	0.7851	**0.7537**	0.7965	**0.7548**	**0.6784**	0.8105	0.7816	**0.7407**
	Fisher	0.8890	0.7576	0.6077	0.8076	0.5566	0.3298	0.8903	0.7539	0.6032
	Sum	0.8778	0.7994	0.7142	**0.8400**	0.7229	0.5821	0.8777	0.7970	0.7046
		(S4)			(S5)			(S6)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	0.7028	0.3953	0.1939	0.6864	0.5398	0.4718	**0.8784**	0.7345	0.5799
	Quantile	0.6945	**0.5850**	**0.4308**	0.8213	**0.8025**	**0.7791**	0.8068	**0.7770**	**0.7232**
	Fisher	0.6867	0.3234	0.1142	0.7725	0.4868	0.2892	0.8671	0.6939	0.5033
	Sum	**0.7468**	0.5241	0.3013	**0.8550**	0.7569	0.6297	0.8675	0.7708	0.6634

in all cases the AUC is less than (or approximately) 0.5. These results are interesting to see the amount of excess conservativeness in the conformal data contamination tests, noting that an AUC of (approximately) 0.5 means there is no excess conservativeness. Generally, the Quantile test shows the least amount of excess conservativeness with AUC approximately 0.5 even when $\pi = \pi_{\text{th}} = 0.2$ for most scenarios. Any excess conservativeness can in these results be attributed to the inequalities in Eqs. (A.1) and (A.3), in which an outlier with a conformal p-value above $1/(n - \ell + 1)$ contributes to a lack of tightness in the bound.

D.2 Multiple testing with conformal data contamination p-values

Now, we will investigate the TDR and FDR when doing multiple conformal data contamination tests. The TDR when using Storey's BH procedure at significance level $\alpha = 0.1$ and for Storey's hyperparameter $\gamma = 0.5$ is shown in Table A6. Mostly, the results follow that observed in Table A4 for the marginal power, i.e., the Storey test statistic tends to be the most powerful, but for some cases for larger data contamination thresholds, the Quantile test statistic is preferable. This tells us that the marginal power of the test is a good indication of the power of multiple testing.

Table A7: FDR computed from 2000 simulations at significance level $\alpha = 0.1$ using the BH procedure. For $K_0 = K/2$ of the data agents $\pi_k = \pi_{\text{th}}$ while for the remaining data agents $\pi_k = 1 - \pi_{\text{th}}$. Theoretically, the FDR is upper bounded by $\alpha K_0/K = 0.05$.

		(S1)			(S2)			(S3)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	0.0363	0.0246	0.0251	0.0379	0.0197	0.0115	0.0414	0.0207	0.0234
	Quantile	0.0410	0.0391	0.0361	0.0447	0.0401	0.0324	0.0416	0.0403	0.0374
	Fisher	0.0164	0.0093	0.0067	0.0168	0.0066	0.0029	0.0193	0.0095	0.0052
	Sum	0.0309	0.0228	0.0157	0.0329	0.0171	0.0116	0.0318	0.0230	0.0166
		(S4)			(S5)			(S6)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	0.0416	0.0141	0.0050	0.0390	0.0287	0.0259	0.0418	0.0222	0.0145
	Quantile	0.0447	0.0309	0.0184	0.0454	0.0424	0.0369	0.0472	0.0422	0.0317
	Fisher	0.0182	0.0026	0.0008	0.0170	0.0064	0.0022	0.0179	0.0081	0.0022
	Sum	0.0332	0.0141	0.0055	0.0313	0.0206	0.0118	0.0344	0.0225	0.0105

Table A8: FDR computed from 2000 simulations at significance level $\alpha = 0.1$ using Storey's BH procedure with parameter $\gamma = 0.5$. For $K_0 = K/2$ of the data agents $\pi_k = \pi_{\text{th}}$ while for the remaining data agents $\pi_k = 1 - \pi_{\text{th}}$.

		(S1)			(S2)			(S3)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	0.1064	0.0708	0.0489	0.1105	0.0513	0.0230	0.1159	0.0639	0.0461
	Quantile	0.1120	0.1074	0.1050	0.1198	0.1051	0.0984	0.1128	0.1089	0.0975
	Fisher	0.0641	0.0281	0.0153	0.0643	0.0191	0.0074	0.0653	0.0293	0.0164
	Sum	0.0890	0.0658	0.0458	0.0965	0.0589	0.0347	0.0940	0.0632	0.0483
		(S4)			(S5)			(S6)		
		$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$	$\pi_{\text{th}}=0$	$\pi_{\text{th}}=0.1$	$\pi_{\text{th}}=0.2$
(C8)	Storey	0.1151	0.0341	0.0100	0.1143	0.0736	0.0575	0.1093	0.0625	0.0294
	Quantile	0.1169	0.0932	0.0641	0.1223	0.1094	0.1065	0.1207	0.1095	0.0969
	Fisher	0.0648	0.0089	0.0018	0.0634	0.0196	0.0059	0.0654	0.0283	0.0067
	Sum	0.0967	0.0424	0.0158	0.0941	0.0632	0.0392	0.0947	0.0694	0.0374

To finally validate Theorem 3.5 on these real datasets, we show in Table A7 the FDR when using the BH procedure at level $\alpha = 0.1$ when at the boundary of the null hypothesis. Specifically, we let $K_0 = K/2$ of the data agents have $\pi_k = \pi_{\text{th}}$ and the others have $\pi_k = 1 - \pi_{\text{th}}$ (this maintains a balance of approximately half the data points being outliers). Hence, here we do not simulate $\pi_k \sim \text{Uniform}([0, 1])$ as we do not allow any $\pi_k < \pi_{\text{th}}$ which would inherently cause excess conservativeness. By Theorem 3.5, the FDR should in this case be upper bounded by $\alpha K_0/K = 0.05$, which is also observed in Table A7 for all tested cases. However, we can note that some excess conservativeness arises when applying the BH procedure to the conformal data contamination p-values. Once again, the Quantile test is generally the least conservative and in some of the tested scenarios maintain an FDR close to the 0.05 level even for $\pi_{\text{th}} = 0.2$. On the other hand, the Fisher test tends to be the most conservative out of the four proposed conformal data contamination tests.

Finally, we consider, in Table A8, the FDR when using Storey's BH procedure. Here, you see that the tests are close to the nominal level $\alpha = 0.1$, even for $\pi_{\text{th}} = 0.2$, especially the Quantile tests. Some conservativeness is induced from the multiplicity correction, as was reported above, however, using the adaptive procedure of Storey's BH tends to do the opposite. Recall that we have only shown FDR control with the BH procedure in Theorem 3.5, not with Storey's BH procedure. It will be an interesting line of further work to determine if

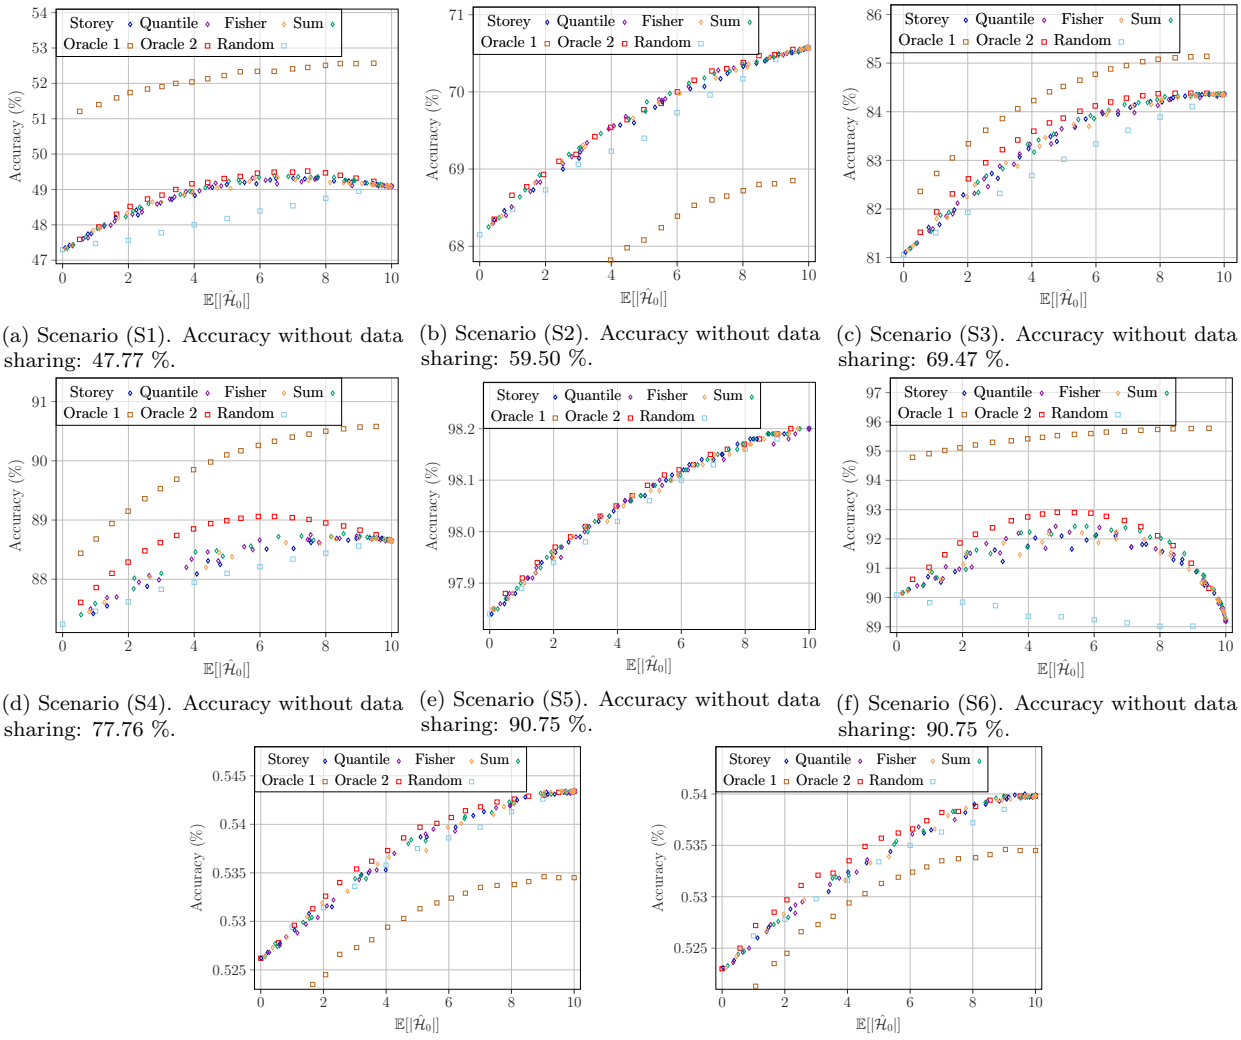

(a) Scenario (S1). Accuracy without data sharing: 47.77 %.

(b) Scenario (S2). Accuracy without data sharing: 59.50 %.

(c) Scenario (S3). Accuracy without data sharing: 69.47 %.

(d) Scenario (S4). Accuracy without data sharing: 77.76 %.

(e) Scenario (S5). Accuracy without data sharing: 90.75 %.

(f) Scenario (S6). Accuracy without data sharing: 90.75 %.

(g) Scenario (S7). Accuracy without data sharing: 46.90 %.

(h) Scenario (S8). Accuracy without data sharing: 46.90 %.

Figure A6: Plot of the model accuracy against the average number of collaborating data agents.

any adaptive BH procedure can maintain FDR control, and also to find a non-trivial lower bound on the FDR when applying the BH procedure to the conformal data contamination p-values.

D.3 Collaborative data sharing

In this section, we numerically evaluate the performance of the proposed data sharing procedure in terms of classification accuracy. This is done for data contamination scenarios (S1)-(S8) with the class-wise score based on PU classification using LR, i.e. (C8), and to solve the classification task we use the well-studied, simple, and robust method in SVC.

We show in Figure A6 the relation between the average number of collaborating data agents, $\mathbb{E}[|\hat{\mathcal{H}}_0|]$, and the average classification accuracy across 500 data simulations for each of the eight data contamination scenarios, with various hyperparameter settings $(\pi_{\text{th}}, \alpha) \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7\} \times \{0.05, 0.2, 0.5, 0.7\}$. No data subset selection is used meaning that all acquired datapoints are used for model fitting, since the focus of this work is on the proposed conformal data contamination tests.

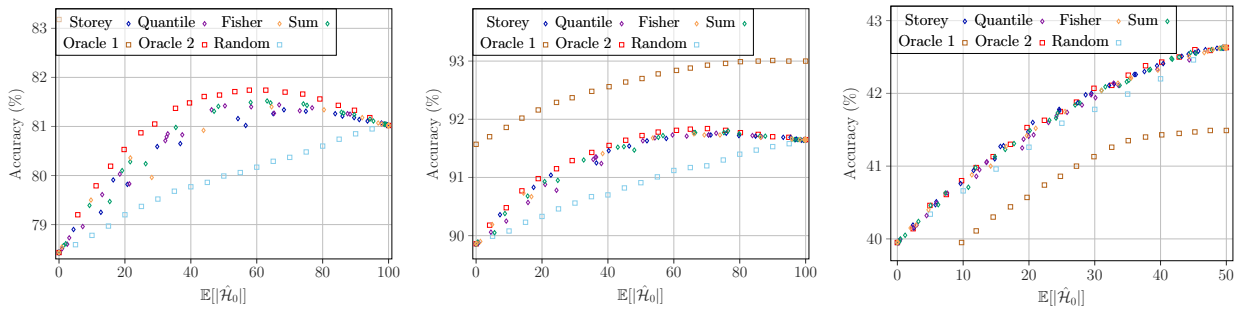

(a) FEMNIST. Accuracy without data sharing: 65.80 %.

(b) MNIST label noise. Accuracy without data sharing: 84.05 %.

(c) CIFAR-10C brightness. Accuracy without data sharing: 32.84 %.

Figure A7: Plot of the model accuracy against the average number of collaborating data agents.

We observe that in scenarios (S1), (S3), (S4), and (S6) the oracles and the proposed methods outperform the *random* baseline meaning that in those scenarios a substantial boost in accuracy can be achieved by carefully selecting data agents for collaboration. Also, in these scenarios *Oracle 1* is clearly outperforming *Oracle 2* showing that there is further room for improvement by incorporating data subset selection. Meanwhile, in scenarios (S2), (S5), (S7), and (S8), i.e., the scenarios with optical coherence tomography images, feature noise on MNIST, and feature noise on CIFAR-10, it appears that the contaminated data is (nearly) as useful for subsequent classification model training as the inliers. In all the considered scenarios, the proposed data sharing procedures outperform the case of no data sharing, emphasizing the value of collaboration. Moreover, the proposed data sharing procedures reaches the performance of *oracle 2* in most cases, showing the efficiency of the proposed methodology. A notable exception is scenario (S4), i.e., the FEMNIST data contamination scenario, and we also observed in Table A3 that outlier detection was difficult in this scenario, hence, showing that in this case it would be beneficial to further improve the conformal score.

These results are limited to small scale problems with small dataset sizes, few collaborators, and few image classes. In the following, we study numerically problems of larger scale.

D.4 Scalability to larger problems

Here, we explore the performance of the data acquisition protocol for problems of larger scale. We consider three scenarios: full MNIST with label noise, FEMNIST with lower- and uppercase handwritten letters *a* through *j*, and full CIFAR-10 mixed with CIFAR-10C contaminated with brightness. All three scenarios have 10 classes, and we set $\ell = 300$, $n - \ell = 250$, and $m = 50$. Further, for MNIST and FEMNIST we set the number of data agents to be $K = 100$, while for the CIFAR-10 scenario $K = 50$ due to limited amount of total data. As previously, we use SVC as the classifier, f^*. For the MNIST and FEMNIST data we use the conformal score (C8) and for CIFAR-10 we use conformal score (C11). Results are only shown for the proposed methods with hyperparameters $(\pi_{\text{th}}, \alpha) \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7\} \times \{0.01, 0.05, 0.2, 0.5, 0.7\}$ when averaged across 20 scenario simulations for MNIST and FEMNIST scenarios, and 100 for the CIFAR-10 scenario. Note that the variability in accuracy across data simulations is reduced in large scale problems.

The results are shows in Figure A7. In the MNIST and FEMNIST scenarios, substantial improvements in accuracy is achieved as compared to the *random* baseline, and in most cases the proposed methods are capable of (nearly) reaching the performance of *oracle 2*, as in the small scale problems previously considered. Generally, the performance gains of the proposed protocols, as compared to the *random* baseline, tends to be more pronounced here than in the small scale problems. This verifies that the proposed data acquisition protocol scales well to problems of larger scale which is an essential property for practical deployment. The accuracy improvements over the *random* baseline is lesser in the CIFAR-10 scenario as we encounter a case where out-of-distribution data still proves helpful to improve learning.

Table A9: Mean (standard deviation) of accuracy (%) across 500 data simulations with data-driven hyperparameter selection.

	Oracle 2	No sharing	Random	Storey	Quantile	Fisher	Sum
(S1)	48.64 (3.00)	47.77 (2.90)	47.96 (2.98)	48.55 (3.02)	48.77 (3.07)	48.56 (3.00)	48.25 (3.06)
(S2)	69.49 (3.22)	59.50 (5.42)	69.35 (3.27)	69.48 (3.34)	69.20 (3.33)	69.36 (3.27)	69.00 (3.34)
(S3)	83.96 (1.59)	69.47 (4.39)	83.75 (1.73)	83.03 (2.37)	83.08 (2.33)	83.09 (2.34)	82.42 (2.27)
(S4)	88.78 (1.47)	77.66 (3.11)	88.22 (1.53)	88.50 (1.56)	88.44 (1.62)	88.46 (1.62)	88.42 (1.65)
(S5)	98.02 (0.28)	90.75 (2.03)	97.97 (0.29)	97.99 (0.30)	98.01 (0.28)	98.01 (0.30)	97.91 (0.31)
(S6)	92.40 (4.01)	90.75 (2.03)	90.86 (5.09)	91.71 (4.58)	91.86 (4.67)	91.93 (4.36)	91.57 (4.87)

D.5 Data-driven hyperparameter selection

For the proposed collaborative data sharing procedure there are a number of hyperparameters, see Section B for an overview. In this section, we will consider a data-driven approach to selecting key hyperparameters. We observed in Figure A6 for some of the data contamination scenarios a concave relation between the accuracy and $\mathbb{E}[|\hat{\mathcal{H}}_0|]$, and so in practice we may be interested in selecting the optimal hyperparameters in terms of accuracy.

The data-driven hyperparameter selection we propose here can be summarized as: (i) order the K data agents in terms of contamination and determine $\hat{\mathcal{H}}_0$; (ii) fit the classification model on the data $\{Z_i\}_{i=1}^{\ell} \bigcup \{Z_i^k : i \in [m], k \in \hat{\mathcal{H}}_0\}$; (iii) evaluate the accuracy on the data $\{Z_i\}_{i=\ell+1}^{n}$ denoted $L_{\text{val}}(\pi_{\text{th}}, \alpha)$ noting that $\hat{\mathcal{H}}_0$ depends on the hyperparameters π_{th} and α; (iv) set the hyperparameter as $(\hat{\pi}_{\text{th}}, \hat{\alpha}) = \arg\max_{(\pi_{\text{th}},\alpha)\in G} L_{\text{val}}(\pi_{\text{th}}, \alpha)$ where G is a grid of hyperparameter values.

We compare with the *random* and *oracle 2* baselines: for the *random* baseline the hyperparameter is K_{budget}, and for baseline *oracle 2* the hyperparameter is π_{th}.

In Table A9, we show the mean and standard deviation of the accuracy (%) across 500 data simulations when using the data-driven hyperparameter selection. In all cases, we consider the settings as described in Section D and use the conformal score (C9). First, we notice that for scenarios (S2)-(S5) all the methods are close in terms of accuracy, since in these scenarios the general rule is that more data is almost always better, as has also been noted above. On the other hand, for scenarios (S1) and (S6) notable gains in accuracy can be achieved with the proposed method. We further observe that we are not quite able to reach the peak accuracies observed in Figure 3 with the proposed methods, however, it is quite close. For this reason, we deem that the data-driven hyperparameter selection paves the way for a practical deployment of the procedure, and note that the optimization of other hyperparameters could also be done with the same approach. Notice also that for the *random* baseline, the mean accuracy surpasses the highest accuracy observed in Figure 3, due to the implicit dependency on the contamination factors in choosing K_{budget}.

This data-driven hyperparameter selection still lacks attention to the cost of collaboration. For instance, in scenarios (S2)-(S5), where more data is almost always better, the optimal choice is often to collaborate with all the data agents. In a specific application, the loss function could be defined to take the cost of collaboration into account.

E Implementation details of conformal scores

Here we provide additional implementation details for the conformal score functions considered in Section D.

Scores (C1)-(C4): The scores (C1)-(C4) are based on the one-class classification methods, one-class support vector machines and isolation forest, and uses the standard implementation in `scikit-learn`. For scores (C1) and (C3), the one-class classifier is fitted to the data $X_1, \ldots, X_\ell$. The calibration and test scores are then evaluated using the scoring function of the fitted model. For the class-wise scores (C2) and (C4), $|\mathcal{Y}|$ one-class classifiers are fitted, one to each dataset $X_{\text{train}}^{Y=i}$ for $i \in \mathcal{Y}$, and again calibration and test scores are evaluated using the score function.

Scores (C5)-(C6): The scores (C5) and (C6) are also one-class classification methods, but autoencoder based. The autoencoder consists of an encoder neural network and a decoder neural network, both having a single layer. The dimension of the latent is 64 with a `ReLu` activation in the hidden layer, and a `sigmoid` activation in the output layer. The autoencoder is trained using 100 epochs with the standard implementation of the `adam` algorithm in `Keras`, employing the mean squared error loss. The calibration and test scores are evaluated as the reconstruction error using the fitted autoencoder. For the score (C5), a single autoencoder is trained on the data $X_1, \ldots, X_\ell$, while for the class-wise score (C6), $|\mathcal{Y}|$ autoencoders are fitted, one to each dataset $X_{\text{train}}^{Y=i}$ for $i \in \mathcal{Y}$.

Scores (C7)-(C14): The scores (C7)-(C14) are based on positive-unlabeled classification (Marandon et al., 2024) using either logistic regression (C7)-(C8), support vector machines (C9)-(C10), logistic regression on principal components explaining at least 80 % of the variance, or a residual network. Logistic regression and support vector machines implementation uses the standard in `scikit-learn` with balanced class weight. Dimensionality reduction to principal components is also implemented using the standard in `scikit-learn`, and is fitted to the training data, i.e., $X_1, \ldots, X_\ell$.

The residual network consists of two convolutional layers, followed by a residual block with two convolutional layers, and finally a fully connected output layer. The first convolutional layer has 32 output channels, uses a kernel of size 3×3, stride of 2, and padding of 1, followed by batch normalization with 0.9 momentum and `ReLu` activation. The second convolutional layer has 64 output channels, uses a kernel of size 3×3, stride of 2, and padding of 1, again followed by batch normalization, `ReLu` activation, and max pooling with kernel size of 2×2 and a stride of 2. The residual block uses a kernel of size 3×3 and no stride. This is followed by average pooling into a single scalar for each of the 64 channels. Finally, the output layer maps with a fully connected layer to a 2-dimensional output with a `sigmoid` activation. We use the cross entropy loss and L_2 weight decay with parameter 0.0001. The model is optimized using the `adam` algorithm with learning rate of 0.0001. The batch size is set to 32 and we use 10 training epochs. The residual network is implemented in `PyTorch`.

For scores (C7), (C9), (C11), and (C13), the binary classifier is fitted to the dataset

$$\{(X_j, 1)\}_{j=1}^{\ell} \bigcup \{(X_j, 0)\}_{j=\ell+1}^{n} \bigcup \{(X_j^k, 0) : j \in [m_{r0}], k \in [K]\}.$$

The calibration and test scores are then evaluated as the estimated probability of belonging to the positive class 1. For the class-wise scores (C8), (C10), (C12), and (C14), $|\mathcal{Y}|$ binary classifiers are fitted, one to each dataset

$$\{(X_j, 1) : Y_j = i, j \in [\ell]\} \bigcup \{(X_j, 0) : Y_j = i, j \in \ell + [n - \ell]\} \bigcup \{(X_j^k, 0) : Y_j^k = i, j \in [m_{r0}], k \in [K]\}.$$

for $i \in [|\mathcal{Y}|]$. Again calibration and test scores are evaluated as the estimated probability of belonging to the positive class 1.

References

Andrea Acevedo, Anna Merino, Santiago Alférez, Ángel Molina, Laura Boldú, and José Rodellar. A dataset of microscopic peripheral blood cell images for development of automatic recognition systems. *Data in Brief*, 30:105474, 2020. ISSN 2352-3409. doi: 10.1016/j.dib.2020.105474.

Stephen Bates, Emmanuel Candès, Lihua Lei, Yaniv Romano, and Matteo Sesia. Testing for outliers with conformal p-values. *The Annals of Statistics*, 51(1):149 – 178, 2023. doi: 10.1214/22-AOS2244.

Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, 29(4):1165 – 1188, 2001. doi: 10.1214/aos/1013699998.

Yoav Benjamini, Abba M. Krieger, and Daniel Yekutieli. Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507, 09 2006. ISSN 0006-3444. doi: 10.1093/biomet/93.3.491.

Christophe A. N. Biscio, Adrien Mazoyer, and Martin V. Vejling. Conformal novelty detection for replicate point patterns with FDR or FWER control. *Spatial Statistics*, 69, 2025. ISSN 2211-6753. doi: 10.1016/j.spasta.2025.100924.

Alexandra Bodzas, Pavel Kodytek, and Jan Zidek. A high-resolution large-scale dataset of pathological and normal white blood cells. *Data in Brief*, 10, 2023. doi: 10.1038/s41597-023-02378-7.

Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. Leaf: A benchmark for federated settings, 2019. URL `https://arxiv.org/abs/1812.01097`.

Emma Dugas, Jared, Jorge, and Will Cukierski. Diabetic retinopathy detection, 2015. URL `https://kaggle.com/competitions/diabetic-retinopathy-detection`.

Varun Gulshan, Lily Peng, Marc Coram, Martin C. Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, Ramasamy Kim, Rajiv Raman, Philip C. Nelson, Jessica L. Mega, and Dale R. Webster. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Journal of American Medical Association*, 316(22):2402–2410, 12 2016. ISSN 0098-7484. doi: 10.1001/jama.2016.17216.

Simon Hawkins, Hongxing He, Graham Williams, and Rohan Baxter. Outlier detection using replicator neural networks. In *Data Warehousing and Knowledge Discovery*, pp. 170–180. Springer Berlin Heidelberg, 2002. ISBN 978-3-540-46145-6.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, Proceedings of Machine Learning Research. PMLR, 2019.

Daniel S. Kermany, Michael Goldbaum, Wenjia Cai, Carolina C.S. Valentim, Huiying Liang, Sally L. Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, Justin Dong, Made K. Prasadha, Jacqueline Pei, Magdalene Y.L. Ting, Jie Zhu, Christina Li, Sierra Hewett, Jason Dong, Ian Ziyar, Alexander Shi, Runze Zhang, Lianghong Zheng, Rui Hou, William Shi, Xin Fu, Yaou Duan, Viet A.N. Huu, Cindy Wen, Edward D. Zhang, Charlotte L. Zhang, Oulan Li, Xiaobo Wang, Michael A. Singer, Xiaodong Sun, Jie Xu, Ali Tafreshi, M. Anthony Lewis, Huimin Xia, and Kang Zhang. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131.e9, 2018. ISSN 0092-8674. doi: 10.1016/j.cell.2018.02.010.

Konstantinos Konstantinou, Tomáš Mrkvička, and Mari Myllymäki. The power of visualizing distributional differences: formal graphical n-sample tests. *Computational Statistics*, 2024. doi: 10.1007/s00180-024-01569-z.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pp. 413–422, 2008. doi: 10.1109/ICDM.2008.17.

Ruhan Liu, Xiangning Wang, Qiang Wu, Ling Dai, Xi Fang, Tao Yan, Jaemin Son, Shiqi Tang, Jiang Li, Zijian Gao, Adrian Galdran, J.M. Poorneshwaran, Hao Liu, Jie Wang, Yerui Chen, Prasanna Porwal, Gavin Siew Wei Tan, Xiaokang Yang, Chao Dai, Haitao Song, Mingang Chen, Huating Li, Weiping Jia, Dinggang Shen, Bin Sheng, and Ping Zhang. DeepDRiD: Diabetic retinopathy – grading and image quality estimation challenge. *Patterns*, 3(6):100512, 2022. ISSN 2666-3899. doi: https://doi.org/10.1016/j.patter.2022.100512.

Ariane Marandon, Lihua Lei, David Mary, and Etienne Roquain. Adaptive novelty detection with false discovery rate guarantee. *The Annals of Statistics*, 52(1):157 – 183, 2024. doi: 10.1214/23-AOS2338.

David Mary and Etienne Roquain. Semi-supervised multiple testing. *Electronic Journal of Statistics*, 16(2): 4926 – 4981, 2022. doi: 10.1214/22-EJS2050.

Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001. doi: 10.1162/089976601750264965.

Saman Sotoudeh-Paima, Ata Jodeiri, Fedra Hajizadeh, and Hamid Soltanian-Zadeh. Multi-scale convolutional neural network for automated amd classification using retinal oct images. *Computers in Biology and Medicine*, 144:105368, 2022. ISSN 0010-4825. doi: 10.1016/j.compbiomed.2022.105368.

John D. Storey, Jonathan E. Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(1):187–205, 12 2003. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2004.00439.x.