



ADAGES: Adaptive Aggregation with Stability for Distributed Feature Selection

Yu Gui

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

August 2, 2020

ADAGES: adaptive aggregation with stability for distributed feature selection

Yu Gui*

July 28, 2020

Abstract

In this era of “big” data, not only the large amount of data keeps motivating distributed computing, but concerns on data privacy also put forward the emphasis on distributed learning. To conduct feature selection and to control the false discovery rate in a distributed pattern with multi-*machines* or multi-*institutions*, an efficient aggregation method is necessary. In this paper, we propose an adaptive aggregation method called ADAGES which can be flexibly applied to any machine-wise feature selection method. We will show that our method is capable of controlling the overall FDR with a theoretical foundation while maintaining power as good as the Union aggregation rule in practice.

1 Introduction

In recent decades, the idea of distributed learning and data decentralization has been frequently discussed. On one hand, the notion of distributed learning is motivated by the advanced techniques of data collection and storage which leads to a large amount of accessible data. Distributed storage and parallel computing are put forward to address the concerns, which further requires statistical learning methods in this distributed scenario. On the other hand, statisticians focus on distributed learning since privacy protection is of main interest nowadays. A representative example is the collaborative clinical research among different hospitals on certain diseases, where hospitals will not

*The author finished this paper when he was an undergraduate at the University of Science and Technology of China

share patients’ data for privacy protection. Therefore, statisticians have to deal with certain “encoded” statistics collected from distributed institutions.

Many recent works focusing on different statistical perspectives have contributed to this field. Estimation is the most fundamental topic in statistics, some works adopt the divide and conquer algorithm for distributed estimation and also study the accuracy of estimation under various contexts, among which are (Battey et al., 2015), (?), (Zhao et al., 2014) and (Cai and Wei, 2020). Distributed hypothesis testing is discussed in works such as (Ramdas et al., 2017), (Sreekumar et al., 2018), (Gilani et al., 2019) and is also covered in (Battey et al., 2015) and (Zhao et al., 2014). Specifically, (Su et al., 2015), (Emery and Keich, 2019) and (Nguyen et al., 2020) have studied the aggregated feature selection based on multiple knockoffs. Originated from applications, communication constraints and privacy constraints ought to be taken into consideration, (Zhang and Berger, 1988), (Braverman et al., 2016), (Cai and Wei, 2020) study the tradeoff between communication constraints and estimation accuracy. In addition, many other works contribute to distributed learning theories such as (Garg et al., 2014), (Dobriban and Sheng, 2018), (Jordan et al., 2019) and (Kipnis and Duchi, 2019).

Controlled feature selection. In addition to feature selection methods such as regularized regression (e.g. (?),(?)), controlled feature selection aims to select important features and reduce false selections under some criteria. In this paper, we focus on a fundamental criterion in feature selection: false discovery rate (FDR). The notion of FDR is introduced in (Benjamini and Hochberg, 1995). With the definition of the subset $\mathcal{S} \subset \{1, \dots, d\}$ of relevant features, feature selection is equivalent to recovering \mathcal{S} based on observations. When the estimated set $\hat{\mathcal{S}}$ is produced, the false discoveries can be denoted as $\hat{\mathcal{S}} \cap \mathcal{S}^c$ and *false discovery proportion* (FDP) is defined in the form

$$\text{FDP} = \frac{|\hat{\mathcal{S}} \cap \mathcal{S}^c|}{|\hat{\mathcal{S}}|}. \tag{1.1}$$

The expectation of FDP is called the *false discovery rate* (FDR), i.e. $\text{FDR} = \mathbb{E}[\text{FDP}]$. In addition, power of feature selection illustrates the ability to recover true features and thus is defined as

$$\text{Power} = \mathbb{E}\left[\frac{|\hat{\mathcal{S}} \cap \mathcal{S}|}{|\mathcal{S}|}\right], \tag{1.2}$$

which is the expected number of true discoveries over the total number of true features




Figure 1: Distributed feature selection and aggregation process

$|\mathcal{S}|$.

A series of FDR-based methods originate from the invention of FDR in (Benjamini and Hochberg, 1995) which utilizes the rank of z-scores for selecting important features. Based on this, (Benjamini et al., 2001) relaxes the independence assumption as an extension. Knockoff filter is introduced in (Barber et al., 2015) with exact control of FDR and can be extended in a model-free way in (Candès et al., 2016). Recently, methods based on mirror statistics are put forward under this topic: (Xing et al., 2019) creates Gaussian mirror variables for all features that get rid of the conditional correlation within each mirrored pair; (Dai et al., 2020) utilizes the data splitting and multiple splitting techniques to ensure the recovery of feature importance with stability.

Stability selection. As an improvement to general feature selection methods, the notion of stability selection is introduced by (Meinshausen and Bühlmann, 2008) which conducts subsampling of size $\lfloor n/2 \rfloor$ and identifies the most frequently selected features. The idea is close to a “voting process” where each sub-sample votes for each feature once and it is in line with our belief that important features will stably become outstanding with more votes. The spirit of stability selection later motivates works such as (Shah and Samworth, 2011), (Hofner et al., 2015) and also stimulates our idea of adaptive aggregation in distributed feature selection.

Our contribution. With the belief in the future of data decentralization, in this paper, we consider the topic of distributed feature selection with a controlled error rate. We present a general aggregation method for distributed feature selection called ADAGES (**AD**aptive **AG**grEgation with **S**tability) that can apply for any controlled feature selection method. Without looking into the original datasets, we operate on Boolean variables in

$\{0, 1\}^d$ that is equivalent to subset of features of dimension d . Therefore, there is no complex communication or privacy concern in this context. Unlike (Su et al., 2015), (Emery and Keich, 2019) and (Nguyen et al., 2020) that transfer knockoff statistics for aggregation, ADAGES does not depend on any specific feature selection method and is thus more flexible in application.

Besides, in this paper, we assume the feature selection procedures of all the machines are independent of each other, i.e. as random Boolean vectors, $\widehat{S}_i \perp \widehat{S}_j$ for all $i, j \in [k]$. It is noticeable that in practice, the dependence exists due to the overlap of samples for different machines, e.g. the common patients for different hospitals. The generalized case to study the dependence is a promising topic for future work.

Outline. We begin with the problem formulation in section 2 and then in section 3, we introduce the detail of ADAGES as an adaptive improvement on empirical rules. In section 4, the main theorem will be established to guarantee the exact control of overall FDR, theoretical proofs of which are in section A. The results of numerical experiments are shown in section 5.

Notations. Suppose the dimension of the n observed features is d , i.e. $\mathbf{x} \in \mathbb{R}^d$. Define $S \subset \{1, \dots, d\}$ is the subset of true features of interest. There are k different machines or institutions contributing to the problem and we denote them as M_1, \dots, M_k . For each $i \in \{1, \dots, k\}$, machine M_i produces an estimated subset \widehat{S}_i before aggregation and our goal is to obtain \widehat{S} based on $\{\widehat{S}_1, \dots, \widehat{S}_k\}$. Notation $\widehat{S}_{(c)}$ refers to the subset produced by the aggregation method with threshold c , which will be introduced in section 3. Also, \widehat{S}_I and \widehat{S}_U are aggregated subsets of the Intersection rule and the Union rule respectively.

2 Background

In the context of distributed learning, imagine there is a central machine (the yellow one in Figure 1) and k machines $\{M_i : i = 1, \dots, k\}$ which can be k hospitals or servers. In the current task, the dataset of interest is distributed over all k machines due to concerns of privacy or distance and assume the i th machine deals with a sub-dataset D_i with n_i observations. All the machines share the same set of features in the same task, i.e. $\{X_j : j = 1, \dots, d\}$ and they focus on FDR control with the universal pre-defined level of $q \in (0, 1)$. Suppose the selection result for the i th machine is \widehat{S}_i . We should note that the feature selection method adopted for each machine can be arbitrary and the only requirement is that the method should be capable of exact FDR control. With

our adaptive aggregation with stability, we produce the final selection result $\widehat{\mathcal{S}}$ based on controlled selections $\{\widehat{\mathcal{S}}_i : i = 1, \dots, k\}$. For each machine $M_i, i = 1, \dots, k$, we define $\text{FDR}_i = \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_i|} \right]$.

2.1 Empirical aggregation methods for distributed feature selection

First, three empirical aggregation methods are introduced and we will later cover them as special cases in a generalized family. Define $z_j^{(i)} = \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}}$ for each feature, then $\widehat{\mathcal{S}}_i$ is equivalent to an indicator vector $\mathbf{z}^{(i)} = (z_1^{(i)}, \dots, z_d^{(i)})^\top$ and aggregation algorithms can be viewed as operation rules for Boolean variables. Also, in the sense of privacy protection, the selected subset $\widehat{\mathcal{S}}_i$ as the statistics with less sensitive information can be publicly transferred to the “center machine” for aggregation. Among aggregation methods, union and intersection of sets are usually adopted empirically. As the simplest rule similar to the OR rule in Boolean operation, we obtain the Union rule

$$\widehat{\mathcal{S}}_U = \bigcup_{i=1}^k \widehat{\mathcal{S}}_i. \quad (2.1)$$

Also, the intersection of all selected subsets produces the Intersection rule:

$$\widehat{\mathcal{S}}_I = \bigcap_{i=1}^k \widehat{\mathcal{S}}_i. \quad (2.2)$$

The Union rule is not strict, thus requires the stricter FDR control for each machine. It indicates that if each machine has FDR control at q , then the overall FDR may far exceeds the pre-defined level. On the other hand, the Intersection rule is far more stricter and will result in the loss of power in aggregation. The phenomenon is illustrated in the left plot of Figure 2. We will show that these two rules will have a more general representation and are thus included in a family of threshold-based aggregation rules.

2.2 Generalized threshold-based aggregation

As an extension to the operation of Boolean variables, we first define

$$m_j = \sum_{i=1}^k \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}}, \quad j = 1, \dots, d. \quad (2.3)$$

Then the *threshold-based rule* is conducted as

$$\widehat{\mathcal{S}}_{(c)} = \{j \in [d] : m_j \geq c\} \quad (2.4)$$

for an integer c .

Remark 2.1. We should notice that *the Union rule* is a special case of *the threshold-based rule* with $\widehat{\mathcal{S}}_U = \widehat{\mathcal{S}}_{(c=1)}$. And for *the Intersection rule*, $\widehat{\mathcal{S}}_I = \bigcap_{i=1}^k \widehat{\mathcal{S}}_i = \widehat{\mathcal{S}}_{(c=k)}$.

Lying between the Intersection and the Union rules, the threshold $c = \lceil (k + 1) / 2 \rceil$ can be adopted as a mild rule and we call it “median-aggregation”. However, we rarely have prior information to determine a universal threshold c and the suitable threshold may also vary in different cases. Therefore, we introduce ADAGES, the adaptive aggregation method in the following section.

3 Adaptive aggregation for distributed feature selection

Based on the definition of $\widehat{\mathcal{S}}_{(c)}$, $\widehat{\mathcal{S}}_{(c_1)} \subseteq \widehat{\mathcal{S}}_{(c_2)}$ for any $c_1 \geq c_2$, thus $|\widehat{\mathcal{S}}_{(c)}|$ is a decreasing function of c . Further, adaptive information aggregation from k machines utilizes the data-driven threshold which is determined conditionally on $\{\widehat{\mathcal{S}}_i, i = 1, \dots, k\}$, thus it is meaningful to look into the behavior of $\widehat{\mathcal{S}}_{(c)}|(\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k)$. Denote $\bar{s} = \frac{1}{k} \sum_{i=1}^k |\widehat{\mathcal{S}}_i|$ and $M = \max_{i=1}^k |\widehat{\mathcal{S}}_i|$.

3.1 Candidate region for threshold

Restrictions on the size of $\widehat{\mathcal{S}}_{(c)}$ is one traditional way to regularize model complexity, and in the first step, we determine the candidate region for threshold c by restricting the model complexity measure $|\widehat{\mathcal{S}}_{(c)}|$. In the contrast to the usual upper bounds for model complexity, we use the mean $\bar{s} = \frac{1}{k} \sum_{i=1}^k |\widehat{\mathcal{S}}_i|$ as a lower bound for $|\widehat{\mathcal{S}}_{(c)}|$, which is in line with the purpose of power maintenance in multiple testing.

We define c_0 as an upper bound as

$$c_0 = \max \left\{ c : |\widehat{\mathcal{S}}_{(c)}| \geq \bar{s} \right\} \quad (3.1)$$

and it is trivial that $c_0 \geq 1$ since

$$|\widehat{\mathcal{S}}_U| = |\widehat{\mathcal{S}}_{(c=1)}| \geq \max \left\{ |\widehat{\mathcal{S}}_i| : i = 1, \dots, k \right\} \geq \bar{s}. \quad (3.2)$$




Figure 2: Left: Empirical power and FDP against threshold c ; Right: Subset size $|\widehat{\mathcal{S}}_{(c)}|$ and modified-ratio $20 * \log(\eta_c)$ against threshold c . ($n = 1000, d = 50, k = 10$)

Therefore, we can choose any integer $c \leq c_0$ as a mild threshold for aggregation, but in the meanwhile, a threshold ought to be chosen to balance the tradeoff between false discovery rate and power.

3.2 Choice of threshold for recovery accuracy

Besides, to improve the tradeoff between FDR and Power, we adopt the following rule emphasizing stable recovery. With c_0 as an upper bound, smaller threshold leads to higher selection power as well as more false discoveries.

Complexity ratio. First, we consider the complexity ratio

$$\eta_c = \begin{cases} \frac{|\widehat{\mathcal{S}}_{(c)}|}{|\widehat{\mathcal{S}}_{(c+1)}|}, & |\widehat{\mathcal{S}}_{(c+1)}| > 0, \\ \infty, & |\widehat{\mathcal{S}}_{(c+1)}| = 0, \end{cases} \quad (3.3)$$

for thresholds decreasing from c_0 and the minimum of complexity ratio is a sign of stable and accurate recovery. Then, the adaptive threshold c^* for aggregation can be chosen by

$$c^* = \operatorname{argmin} \{ \eta_c : 1 \leq c \leq c_0, \}. \quad (3.4)$$

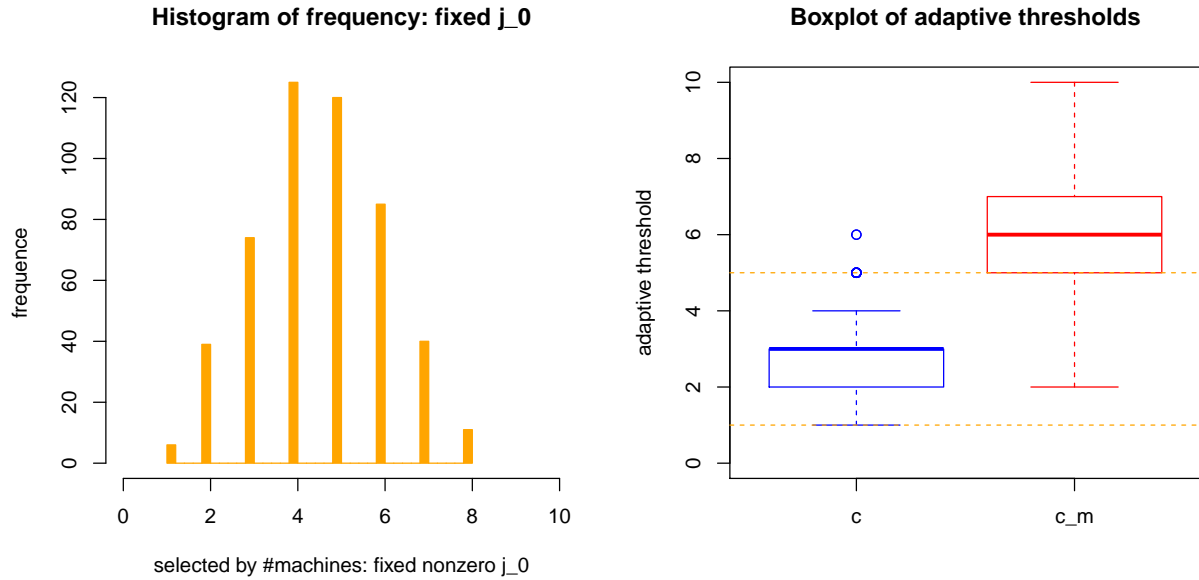


Figure 3: Left: Distribution of appearance in the selected set of $k = 10$ machines for a fixed nonzero feature j_0 over 500 trials; Right: Boxplot of adaptive threshold chosen by ADAGES over 500 trials, $c = c^* = \operatorname{argmin}_{1 \leq c \leq c_0} \eta_c$ and modified threshold $c_m = \tilde{c} = \operatorname{argmin}_{1 \leq c \leq c_0} c|\widehat{\mathcal{S}}_{(c)}|$; Dash lines: $c_U = 1$ and $c_M = \lfloor (k+1)/2 \rfloor$ ($d = 50, n = 1000, \#\text{nonzeros} = 10$).

In practice, to avoid infinite values, we can also use a surrogate $(1 + |\widehat{\mathcal{S}}_{(c)}|)/(1 + |\widehat{\mathcal{S}}_{(c+1)}|)$. As is shown in a simple example in the right plot of Figure 2 with true $|\mathcal{S}| = 20$, threshold with the minimum ratio η_c produces a more stable recovery of the true \mathcal{S} , and in this figure we adopt a modified form $20 \times \log(\eta_c)$ to represent the magnitude of ratio.

Remark 3.1. To illustrate the complexity ratio, the idea is similar to the eigenvalue ratio in PCA for determining the number of meaningful eigen-components. We can also consider a toy example where $|\mathcal{S} \cap \widehat{\mathcal{S}}| \sim B(k, p)$ with $p = \mathbb{P}(j \in \mathcal{S} : j \in \widehat{\mathcal{S}}_i)$. In this case, minimizing the ratio η_c approximately produces the mode of Bernoulli distribution that recovers the threshold in line with the most likely frequency for important features.

Remark 3.2 (Threshold-complexity tradeoff). It is noticeable that another rule with theoretical intuition for choosing the threshold is given by

$$\tilde{c} = \operatorname{argmin}_{1 \leq c \leq c_0} c|\widehat{\mathcal{S}}_{(c)}|, \quad (3.5)$$

which explicitly focuses on the tradeoff between the magnitude of threshold and the size of selected subset. As we will show in Lemma 4.1, the power shrinkage term $\left(\frac{c|\widehat{\mathcal{S}}_{(c)}|}{k|\mathcal{S}|} \cdot \text{FDP}\right)$

plays the leading role in the lower bound for the true positive proportion. Then, for FDP at a certain level, minimizing the product $c|\widehat{\mathcal{S}}_{(c)}|$ is equivalent to maximizing the true positive proportion.

Details of numerical simulations will be discussed in section 5 and the implementation of adaptive aggregation based on the complexity ratio is shown in the Algorithm 1. Aggregated feature selection is an initial case dealing with binary variables. It is more exciting to extend this threshold-based aggregation method to estimation and inference based on communication of more informative statistics, and we leave this for future work.

Algorithm 1 ADAGES: adaptive aggregation with stability for distributed feature selection

- 1: **Input** $\{\widehat{\mathcal{S}}_i : i = 1, \dots, k\}$: $\widehat{\mathcal{S}}_i \subset [d]$ is the selected subset for the i th machine
 - 2: **Output** $\widehat{\mathcal{S}} = \widehat{\mathcal{S}}_{(c^*)}$ as an estimation for \mathcal{S}
 - 3: Calculate $m_j = \sum_{i=1}^k \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}}, \forall j \in \{1, \dots, d\}$
 - 4: Calculate $\bar{s} = \frac{1}{k} \sum_{i=1}^k |\widehat{\mathcal{S}}_i|$
 - 5: **for** c in $\{1, \dots, k\}$ **do**
 - 6: $\widehat{\mathcal{S}}_{(c)} = \{j \in [d] : m_j \geq c\}$
 - 7: Calculate the complexity ratio $\eta_c = \frac{|\widehat{\mathcal{S}}_{(c)}|+1}{|\widehat{\mathcal{S}}_{(c+1)}|+1}, c \leq k-1; \eta_k = \infty$
 - 8: **end for**
 - 9: Determine $c_0 = \max \{c : |\widehat{\mathcal{S}}_{(c)}| \geq \bar{s}\}$
 - 10: Produce adaptive threshold $c^* = \operatorname{argmin} \{\eta_c : 1 \leq c \leq c_0\}$
 - 11: **return** $\widehat{\mathcal{S}} = \widehat{\mathcal{S}}_{(c^*)} = \{j \in [d] : m_j \geq c^*\}$
-

4 Main result

In this section, we will show the theoretical properties of ADAGES for adaptive aggregation in the scenario of distributed feature selection. First, we obtain the control of overall false discovery rate in theorem 4.1; besides, we establish the connection of overall power and machine-wise power: theorem 4.2 shows the simultaneous control of FDR and a power shrinkage term and theorem 4.3 compares the power of ADAGES with the “optimal” power produced by the Union rule.

4.1 Distributed FDR control

Based on the adaptive threshold for aggregation, the ADAGES produces exact control of the false discovery rate.

Theorem 4.1. *For a pre-defined level $q \in (0, 1)$, suppose machine-wise $\text{FDR}_i \leq q$ for $i = 1, \dots, k$ and $\lambda \geq \max_{1 \leq i \leq k} \frac{|\widehat{\mathcal{S}}_i|}{c^*} \sum_{j=1}^k \frac{1}{|\widehat{\mathcal{S}}_j|}$. Then, ADAGES with $c^* \in [1, c_0] \cap \mathbb{Z}$ produces*

$$\text{FDR}_{(c^*)} = \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_{(c^*)} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_{(c^*)}|} \right] \leq \lambda q. \quad (4.1)$$

Then, we discuss two special cases with fixed thresholds $c = 1$ and $c = k$ respectively, which may reveal their shortcomings to some extent.

Proposition 4.1 (the Union rule). *For a pre-defined level $q \in (0, 1)$, if machine-wise $\text{FDR}_i \leq q$ for all $i = 1, \dots, k$, the Union rule produces*

$$\text{FDR}_U = \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_U \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_U|} \right] \leq kq. \quad (4.2)$$

More generally, as is pointed out in (Xie and Lederer, 2019), if there is a sequence of pre-defined FDR levels (q_1, \dots, q_k) such that $\text{FDR}_i \leq q_i$ for all $i \in [k]$, then the overall FDR can be exactly controlled at level $q = \sum_{i=1}^k q_i$. If we would like to have overall FDR controlled at level q , it requires that $\sum_{i=1}^k q_i = q$ and a simple case is $q_i = q/k$ for all k machines. Besides, in the case with $c = k$, based on $k|\widehat{\mathcal{S}}_I \cap \mathcal{S}^c| \leq \sum_{i=1}^k |\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|$, we have the following proposition:

Proposition 4.2 (the Intersection rule). *For a pre-defined level $q \in (0, 1)$, if machine-wise $\text{FDR}_i \leq q$ for $i = 1, \dots, k$ and there is a constant $\kappa \geq 1$ such that $\max_{i \in [k]} \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_I|} \leq \kappa$, then the Intersection rule produces*

$$\text{FDR}_I = \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_I \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_I|} \right] \leq \kappa q. \quad (4.3)$$

Comparing the overall FDR bounds, the Union rule as a less strict aggregation rule produces FDR at an expected level as high as kq . Instead, the Intersection rule is the most conservative and has theoretical FDR control at q multiplied by a factor κ . However, with an adaptive threshold, ADAGES summarizes machine-wise information more

efficiently and has the control of overall FDR at level λq . Here, as an illustration, we compare the magnitude of k, λ, κ to show the abilities of FDR control of the three methods. First, if c^*/k has a positive lower bound such that $c^* \geq b \cdot k$ and $\max_{i \in [k]} \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_j|} = O(1)$ for all j , then we obtain $\lambda = o(k)$. Comparison between λ and κ is of more interest, which is summarized in the following proposition.

Proposition 4.3. *Denote the tight bound $\bar{\lambda} = \max_{i \in [k]} \frac{|\widehat{\mathcal{S}}_i|}{c^*} \sum_{j=1}^k \frac{1}{|\widehat{\mathcal{S}}_j|}$ and $\bar{\kappa} = \max_{i \in [k]} \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_j|}$. Then, we have*

$$\frac{\bar{\lambda}}{\bar{\kappa}} = \frac{1}{c^*} \sum_{j=1}^k \frac{|\widehat{\mathcal{S}}_j|}{|\widehat{\mathcal{S}}_j|}. \quad (4.4)$$

Further, if $(1 - \epsilon)c^* < \sum_{j=1}^k \frac{|\widehat{\mathcal{S}}_j|}{|\widehat{\mathcal{S}}_j|} < (1 + \epsilon)c^*$ for any $\epsilon \in (0, 1)$, then $|\frac{\bar{\lambda}}{\bar{\kappa}} - 1| < \epsilon$.

4.2 Power analysis

We also establish a lower bound for the Power based on $\{\text{Power}_i\}$, $i = 1, \dots, k$ as well as the power produced by the Union bound, before which we introduce the basic lemma to establish the connection between overall true positive proportion (TPP) with machine-wise TPP $_i$, $i = 1, \dots, k$.

Lemma 4.1. *Based on the ADAGES algorithm, we obtain*

$$\text{TPP} \geq \frac{1}{k} \sum_{i=1}^k \text{TPP}_i - \frac{c^* |\mathcal{S}^c \cap \widehat{\mathcal{S}}_{(c^*)}|}{k |\mathcal{S}|}. \quad (4.5)$$

The second term $\frac{c^* |\mathcal{S}^c \cap \widehat{\mathcal{S}}_{(c^*)}|}{k |\mathcal{S}|}$ acts as the term of “power shrinkage” and can be connected with FDP in the form:

$$\text{Power shrinkage} = \frac{c^* |\widehat{\mathcal{S}}_{(c^*)}|}{k |\mathcal{S}|} \text{FDP}, \quad (4.6)$$

which involves a tradeoff between c^* and $|\widehat{\mathcal{S}}_{(c^*)}|$. Therefore, with proper restriction on $|\widehat{\mathcal{S}}_{(c^*)}|$, i.e. a proper choice of c^* , we can simultaneously control FDR and the power shrinkage term, which is shown in theorem 4.2.

Theorem 4.2. *Denote Power_i as the selection of for the i th machine. Suppose there exists constant $\gamma \in (0, 1/2)$ such that $|\widehat{\mathcal{S}}_{(c^*)}| \leq (1 + \gamma)|\mathcal{S}|$ and $c^* \leq k/2$. If the overall FDR is*

controlled at level $q \in (0, 1)$, then for a constant $\alpha \leq 3/4$, we have

$$\text{Power} \geq \frac{1}{k} \sum_{i=1}^k \text{Power}_i - \alpha q. \quad (4.7)$$

It is noticeable that power produced by the Union bound is the maximum power one aggregation method can achieve. Denote $\text{diff} = |(\widehat{\mathcal{S}}_U \cap \mathcal{S}) \setminus (\widehat{\mathcal{S}}_U \cap \mathcal{S})| = |(\widehat{\mathcal{S}}_U \cap \mathcal{S})| - |(\widehat{\mathcal{S}} \cap \mathcal{S})|$, with which we obtain the following theorem.

Theorem 4.3. *Suppose we have a uniform lower bound for $\text{Power}_i, i \in [k]$ that $\mathbb{P}(j \in \mathcal{S}, j \in \widehat{\mathcal{S}}_i) \geq \eta_{n,d}$ for $i \in [k], j \in [d]$. If we further have $c^* \leq k/2$, then $\exists \xi \leq 2$ such that*

$$\mathbb{E}[\text{diff}] \leq \xi(1 - \eta_{n,d})|\mathcal{S}|. \quad (4.8)$$

Further, if the selection method has the property that $\eta_{n,d} \rightarrow 1$ as $n, d \rightarrow \infty$, we have $|\text{Power} - \text{Power}_U| \rightarrow 0$ as $n, d \rightarrow \infty$.

5 Numerical simulation

In this section, we study the performance of our adaptive aggregation method by comparisons with the empirical Union, Intersection and median-aggregation rules in simulations. We also compare with the performance of the aggregation method in (Xie and Lederer, 2019), which is a modified version of the Union rule. In numerical simulations, we use model-X knockoffs with second-order construction for each machine which produces exact FDR control, so the method can be named as “model-X knockoffs + ADAGES” to illustrate the procedure. In this case, we are also interested in the comparison between our algorithm-free ADAGES and the knockoff-based aggregation method AKO in (Nguyen et al., 2020). We consider the AKO with BY step-up with theoretical guarantee and use $\gamma = 0.3$ that is adopted in (Nguyen et al., 2020). In experiments, ADAGES refers to our adaptive method with $c^* = \text{argmin} \{ \eta_c : 1 \leq c \leq c_0 \}$ while ADAGES_m is the modified method with threshold $\tilde{c} = \text{argmin}_{1 \leq c \leq c_0} c |\widehat{\mathcal{S}}_{(c)}|$.

A simple linear model is adopted for feature selection:

$$\mathbf{y} = \mathbf{X}\beta + \epsilon, \quad (5.1)$$

where $\mathbf{X} \in \mathbb{R}^{n \times d} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ is the design matrix, where $\Sigma \in \mathbb{R}^{d \times d}$ and $\Sigma_{ls} = \rho^{|l-s|}$ for all $l, s \in [d]$. $\mathbf{y} \in \mathbb{R}^n$ is the vector of n responses and elements in the noise vector ϵ are




Figure 4: Left: FDP against number of machines k ; Right: empirical power against number of machines k ($k \in \{1, 2, 5, 8, 10, 20\}$ with $n = 1000$, $d = 50$ and $s = |\mathcal{S}| = 20$ over 100 trials).

drawn *i.i.d.* from standard Gaussian distribution. Feature importance is revealed in β and $\mathcal{S} = \{j \in [d] : \beta_j \neq 0\}$.

Comparisons are conducted in the following two aspects, in which the repetition number is $r = 100$ and $\rho = 0.25$. We use the criteria of averaged FDP and averaged power as the sample-versions of FDR and power respectively.

5.1 Varying the number of machines k

Since the number of machines is a vital factor in the context of distributed learning, in the first experiment, we vary k among $\{1, 2, 5, 8, 10, 20\}$ with $n = 1000$, $d = 50$ and $s = |\mathcal{S}| = 20$ fixed. Here nonzero elements in true β is drawn *i.i.d.* and uniformly from $\{\pm 2\}$.

From Figure 4, we can see that ADAGES obtains a desirable tradeoff between the averaged FDP and power. As an adaptive aggregation method, ADAGES controls FDP exactly under $q = 0.2$ while achieves power nearly as good as that of the Union rule, which meets the goal of power maintenance for controlled feature selection. For the three empirical methods, although the Union rule maintains power at the highest level, it produces FDP exceeding the pre-defined level $q = 0.2$; the Intersection rule has con-




Figure 5: Left: FDP against dimension d ; Right: empirical power against dimension d ($d \in \{15, 30, 45, 60, 75, 90\}$ with $n = 1000$, $k = 10$ and $s = |\mathcal{S}| = 10$ over 100 trials).

servative control of FDP but results in a serious loss of power in feature selection while the power loss of median-aggregation occurs earlier than ADAGES.

As an improvement for the Union rule on FDR control, the method in (Xie and Lederer, 2019) obtains comparable FDP with the Intersection rule; but since the pre-defined level for each machine becomes $q_i = q/k$, this method will sacrifice power as shown in Figure 4 and is thus limited in application. On the other hand, in this case without ultra-high dimension or strict sparsity, the AKO that transforms more informative “p-values” in aggregation is capable of controlling the averaged FDP around the level κq where $\kappa \leq 3.24$ is given in (Nguyen et al., 2020); power of AKO is lower than other algorithm-free methods when $k < 10$, but remains stable as k increases.

However, the modified ADAGES with $\tilde{c} = \operatorname{argmin}_{1 \leq c \leq c_0} c |\hat{\mathcal{S}}_{(c)}|$ does not produce higher power in experiments since the power shrinkage term indicates the tradeoff between FDP and $c |\hat{\mathcal{S}}_{(c)}|$. Here, FDP is also a function of c which ought not to be ignored in the choice of \tilde{c} .

5.2 Varying dimension d

In the second experiment, we vary dimension d in the set $\{15, 30, 45, 60, 75, 90\}$ while fix model parameters as $n = 1000$, $k = 10$ and $|\mathcal{S}| = 10$. True signal β_j is generated in the same way mentioned above.

In Figure 5, both ADAGES, median-aggregation and the Intersection rule have exact FDP control under $q = 0.2$, but the Union rule suffers from “uncensored” aggregation and cannot control the overall FDP. Partially dependent on the property of the feature selection method adopted for each machine, the power goes down as d increases. But it is noticeable that the Union rule can always achieve the highest power after aggregation and ADAGES shows comparable performance due to the use of an adaptive threshold based on η_c on an interval with an upper bound.

In addition, the aggregation method in (Xie and Lederer, 2019) tends to make null discovery that is $|\widehat{\mathcal{S}}| = 0$ which naturally control FDR at 0 but also have no power. Similar to our findings with varying k , the AKO performs better than empirical aggregation methods as d increases, especially in power; but ADAGES shows better performance in both averaged FDP and empirical power.

6 Discussion

In this paper, we present an adaptive aggregation method called ADAGES for distributed false discovery rate control. Our method utilizes selected subsets from all machines to determine the aggregation threshold and shows better performance in the tradeoff of FDR control and power maintenance compared with empirical aggregation methods. The ADAGES is algorithm-free, which means it can be applied to any machine-wise feature selection method, and is thus more flexible than aggregation rules based on specific statistics produced by each machine-wise method. It is motivating to further study the modified method based on the power shrinkage term, which has theoretical intuition for power maintenance and requires a good estimation of overall FDP.

Besides, as potential extensions, we can adopt this adaptive method with stability in other statistical aspects in distributed learning. Selected subsets are binary vectors consisting of limited but private information and we can further take communication constraints and privacy into consideration, which are left for our future work. More importantly, there is a tradeoff between information communication and selection power, thus it is meaningful to study aggregation methods with machines transferring encoded

but more informative statistics.

As the distributed pattern becomes more common in the statistical community, to promote inter-institutional collaboration, efficient aggregation methods are necessary for distributed computing as well as privacy protection. With the idea of adaptive aggregation, collaboration can adapt to specific scenarios while each institution simply needs to focus on its specific statistical problem, which greatly contributes to the new collaboration mode in data science. However, another direction for future research is to relax the independence assumption among institutions in the learning procedure and to study the influence of inter-institutional dependence in the statistical context.

Implementation of ADAGES with R is available and raw codes can be accessed on https://github.com/yugjerry/ADAGES/blob/master/code_ADAGES.R. Technical proofs are presented in the following sections.

A Technical proofs

In this section, we present the proofs for the main theorems and propositions in this paper.

Proof for Theorem 4.1. With $\text{FDR}_i = \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_i|} \right] \leq q$, observe the overall FDR:

$$\begin{aligned} \text{FDR}_{(c^*)} &= \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_{(c^*)} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_{(c^*)}|} \right] \\ &= \mathbb{E} \left\{ \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_{(c^*)} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_{(c^*)}|} \middle| (\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k) \right] \right\}. \end{aligned} \quad (\text{A.1})$$

First, we have $c^*|\widehat{\mathcal{S}}_{(c^*)}| \leq \sum_{j:m_j \geq c^*} m_j \leq \sum_{j=1}^d m_j = \sum_{i=1}^k |\widehat{\mathcal{S}}_i|$, and similarly for features $j \in \mathcal{S}^c$,

$$c^*|\widehat{\mathcal{S}}_{(c^*)} \cap \mathcal{S}^c| \leq \sum_{j \in \mathcal{S}^c: m_j \geq c^*} m_j \leq \sum_{j \in \mathcal{S}^c} m_j = \sum_{i=1}^k |\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|. \quad (\text{A.2})$$

Therefore, the overall FDR can be linked to the machine-wise FDR's as

$$\text{FDR}_{(c^*)} = \mathbb{E} \left\{ \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_{(c^*)} \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_{(c^*)}|} \middle| (\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k) \right] \right\}$$

$$\leq \mathbb{E} \left\{ \frac{1}{c^*} \sum_{i=1}^k \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_{(c^*)}|} \mid (\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k) \right] \right\}. \quad (\text{A.3})$$

In addition, by definition of c^* in the theorem: $|\widehat{\mathcal{S}}_{(c^*)}| \geq \frac{1}{k} \sum_{i=1}^k |\widehat{\mathcal{S}}_i| \geq \frac{k}{\sum_{i=1}^k 1/|\widehat{\mathcal{S}}_i|}$, we then obtain

$$\begin{aligned} \text{FDR}_{(c^*)} &\leq \mathbb{E} \left\{ \frac{1}{c^*} \sum_{i=1}^k \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_{(c^*)}|} \mid (\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k) \right] \right\} \\ &\leq \mathbb{E} \left\{ \frac{1}{kc^*} \sum_{j=1}^k \sum_{i=1}^k \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_j|} \mid (\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k) \right] \right\} \\ &= \mathbb{E} \left\{ \frac{1}{kc^*} \sum_{1 \leq i, j \leq k} \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_i|} \cdot \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_j|} \mid \{\widehat{\mathcal{S}}_l\}_{l=1}^k \right] \right\} \\ &= \mathbb{E} \left\{ \frac{1}{kc^*} \sum_{i=1}^k \sum_{j=1}^k \mathbb{E} \left[\text{FDP}_i \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_j|} \mid (\widehat{\mathcal{S}}_1, \dots, \widehat{\mathcal{S}}_k) \right] \right\} \\ &\leq \sum_{i=1}^k \frac{1}{kc^*} \text{FDP}_i \cdot \left(\max_{1 \leq i \leq k} |\widehat{\mathcal{S}}_i| \sum_{j=1}^k \frac{1}{|\widehat{\mathcal{S}}_j|} \right) \\ &\leq \sum_{i=1}^k \frac{1}{kc^*} \text{FDP}_i \cdot \lambda c^* \leq \lambda q. \end{aligned} \quad (\text{A.4})$$

Here, λ is a bound that

$$\lambda \geq \max_{1 \leq i \leq k} \frac{|\widehat{\mathcal{S}}_i|}{c^*} \sum_{j=1}^k \frac{1}{|\widehat{\mathcal{S}}_j|}. \quad (\text{A.5})$$

■

Proof for theorem 4.2. We consider the expected number of true discoveries $\mathbb{E}|\widehat{\mathcal{S}} \cap \mathcal{S}|$ and denote $\text{TPP}_i = \frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}|}{|\mathcal{S}|}$, $i \in [k]$ and $\text{TPP} = \frac{|\widehat{\mathcal{S}}_{(c^*)} \cap \mathcal{S}|}{|\mathcal{S}|}$. Then we have

$$\begin{aligned} |\mathcal{S}| \sum_{i=1}^k \text{TPP}_i &= \sum_{i=1}^k |\widehat{\mathcal{S}}_i \cap \mathcal{S}| \\ &= \sum_{i=1}^k \sum_{j \in \mathcal{S}} \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}} = \sum_{j \in \mathcal{S}} m_j \end{aligned}$$

$$\begin{aligned}
&= \sum_{j \in \mathcal{S} \cap \widehat{\mathcal{S}}_{(c^*)}} m_j + \sum_{j \in \mathcal{S}^c \cap \widehat{\mathcal{S}}_{(c^*)}} m_j \\
&\leq k|\mathcal{S} \cap \widehat{\mathcal{S}}_{(c^*)}| + c^*|\mathcal{S}^c \cap \widehat{\mathcal{S}}_{(c^*)}|,
\end{aligned} \tag{A.6}$$

which is equivalent to

$$\text{TPP} \geq \frac{1}{k} \left(\sum_{i=1}^k \text{TPP}_i - c^* \frac{|\mathcal{S}^c \cap \widehat{\mathcal{S}}_{(c^*)}|}{|\mathcal{S}|} \right). \tag{A.7}$$

Based on the assumption with an upper bound on $\widehat{\mathcal{S}}_{(c^*)}$ with $\gamma \in (0, 1/2)$,

$$|\widehat{\mathcal{S}}_{(c^*)}| \leq (1 + \gamma)|\mathcal{S}|, \tag{A.8}$$

we take expectation for the inequality and then obtain

$$\text{Power} \geq \frac{1}{k} \sum_{i=1}^k \text{Power}_i - \alpha q, \tag{A.9}$$

where $\alpha = \frac{c^*}{k}(1 + \gamma) < \frac{3}{4}$. ■

Proof for theorem 4.3. We can write explicitly that

$$\text{diff} = \sum_{j \in \mathcal{S}} \mathbf{1}_{\{0 < m_j < c^*\}}. \tag{A.10}$$

Then, for positive m_j , $\mathbb{E}[\text{diff}] = \sum_{j \in \mathcal{S}} \mathbb{P}(m_j < c^*)$ with

$$\begin{aligned}
\mathbb{P}(m_j < c^*, j \in \mathcal{S}) &\leq \frac{k - \mathbb{E}[m_j | j \in \mathcal{S}]}{k - c^*} \\
&= \frac{k - \mathbb{E}[\sum_{i=1}^k \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}} | j \in \mathcal{S}]}{k - c^*} \\
&\leq \frac{k(1 - \eta_{m,d})}{k - c^*} \\
&\leq \xi(1 - \eta_{m,d}),
\end{aligned} \tag{A.11}$$

where $c^* \leq k/2$ by definition and thus $\xi \leq 2$. ■

Proof for proposition 4.1. With the Union rule, $\widehat{\mathcal{S}}_U = \bigcup_{i=1}^k \widehat{\mathcal{S}}_i$ and thus $\widehat{\mathcal{S}}_i \subset \widehat{\mathcal{S}}_U$ for

all $i = 1, \dots, k$. Since $|\bigcup_{i=1}^k A_i| \leq \sum_{i=1}^k |A_i|$, we apply this fact to $\widehat{\mathcal{S}}_U \cap \mathcal{S}^c = \bigcup_{i=1}^k (\widehat{\mathcal{S}}_i \cap \mathcal{S}^c)$ and consider the overall FDR:

$$\begin{aligned}
\text{FDR} &= \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_U \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_U|} \right] \leq \mathbb{E} \left[\sum_{i=1}^k \frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_U|} \right] \\
&\leq \sum_{i=1}^k \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_i|} \right] \leq \sum_{i=1}^k \text{FDR}_i \\
&\leq \sum_{i=1}^k q_i.
\end{aligned} \tag{A.12}$$

■

Proof for proposition 4.2. With $\widehat{\mathcal{S}}_I = \bigcap_{i=1}^k \widehat{\mathcal{S}}_i$, we have $m_j = \sum_{i=1}^k \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}} = k$ for $j \in \widehat{\mathcal{S}}$. Therefore, we have

$$\begin{aligned}
k|\widehat{\mathcal{S}}_I \cap \mathcal{S}^c| &= \sum_{j \in \widehat{\mathcal{S}}_I \cap \mathcal{S}^c} m_j \leq \sum_{j \in \mathcal{S}^c} m_j \\
&= \sum_{j \in \mathcal{S}^c} \sum_{i=1}^k \mathbf{1}_{\{j \in \widehat{\mathcal{S}}_i\}} = \sum_{i=1}^k |\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|
\end{aligned} \tag{A.13}$$

We then consider the overall FDR,

$$\begin{aligned}
\text{FDR} &= \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_I \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_I|} \right] \leq \mathbb{E} \left[\frac{1}{k} \frac{|\widehat{\mathcal{S}}_I \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_I|} \right] \\
&\leq \mathbb{E} \left[\frac{1}{k} \frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_i|} \cdot \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_I|} \right] \\
&\leq \frac{\kappa}{k} \sum_{i=1}^k \mathbb{E} \left[\frac{|\widehat{\mathcal{S}}_i \cap \mathcal{S}^c|}{|\widehat{\mathcal{S}}_i|} \right] \leq \kappa q.
\end{aligned} \tag{A.14}$$

Here $\kappa \geq 1$ is a constant such that $\max_{i \in [k]} \frac{|\widehat{\mathcal{S}}_i|}{|\widehat{\mathcal{S}}_I|} \leq \kappa$. ■

B Illustration of the aggregation process

In this part, results in four cases are provided to illustrate the connection between overall FDR/power and the machine-wise ones. The k grey bars in each plot are the FDR_i s

or Power_i s for k machines.

From the four cases together with the simulation results in our paper, we can see that ADAGES has a better tradeoff than other methods (the Union rule, the Intersection rule, median aggregation and method in (Xie and Lederer, 2019)). For FDR, all methods except the Union rule produce the exact control whenever machine-wise FDR is controlled at the pre-defined level. The Union rule, however, as is shown in proposition 4.1, is only able to control FDR at a higher level. When k or the dimension d is large, strict aggregation methods will cause the power loss, such as the results of the Intersection rule and method in (Xie and Lederer, 2019). We should note that “strict” refers to strict pre-defined levels for each machine as well as strict aggregation rules. As is shown in the results, ADAGES produces power very close to that of the Union rule, which is the highest power an aggregation method can achieve.




Figure 6: Representation of the aggregation process: barplot of machine-wise FDR(left)/power(right) and aggregation results under different rules ($q = 0.2, k = 5, d = 20, n = 1000, n_i = 200$).




Figure 7: Representation of the aggregation process: barplot of machine-wise FDR(left)/power(right) and aggregation results under different rules ($q = 0.2, k = 5, d = 80, n = 1000, n_i = 200$).




Figure 8: Representation of the aggregation process: barplot of machine-wise FDR(left)/power(right) and aggregation results under different rules ($q = 0.2, k = 10, d = 20, n = 1000, n_i = 100$).

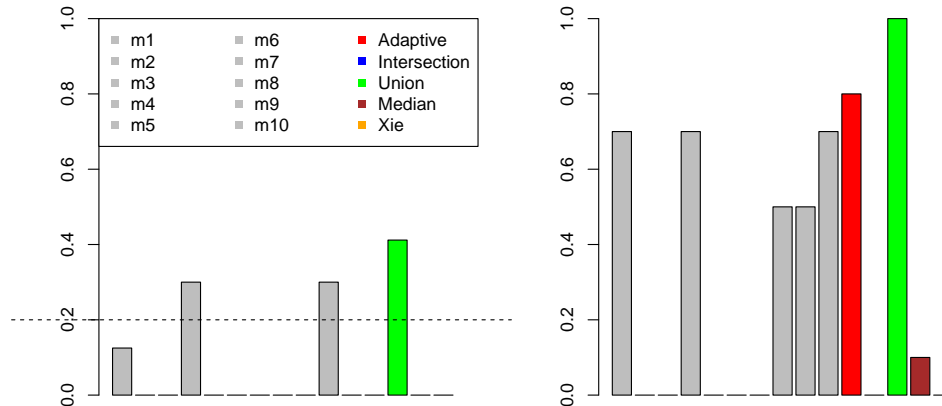


Figure 9: Representation of the aggregation process: barplot of machine-wise FDR(left)/power(right) and aggregation results under different rules ($q = 0.2, k = 10, d = 80, n = 1000, n_i = 100$).

References

- Rina Foygel Barber, Emmanuel J Candès, et al. 2015. Controlling the false discovery rate via knockoffs. *The Annals of Statistics* 43, 5 (2015), 2055–2085.
- Heather Battey, Jianqing Fan, Han Liu, Junwei Lu, and Ziwei Zhu. 2015. Distributed estimation and inference with statistical guarantees. *arXiv preprint arXiv:1509.05457* (2015).
- Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300.
- Yoav Benjamini, Daniel Yekutieli, et al. 2001. The control of the false discovery rate in multiple testing under dependency. *The annals of statistics* 29, 4 (2001), 1165–1188.
- Mark Braverman, Ankit Garg, Tengyu Ma, Huy L. Nguyen, and David P. Woodruff. 2016. Communication Lower Bounds for Statistical Estimation Problems via a Distributed Data Processing Inequality. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing (STOC '16)*. Association for Computing Machinery, New York, NY, USA, 1011–1020. <https://doi.org/10.1145/2897518.2897582>

- T Tony Cai and Hongji Wei. 2020. Distributed Gaussian Mean Estimation under Communication Constraints: Optimal Rates and Communication-Efficient Algorithms. *arXiv preprint arXiv:2001.08877* (2020).
- Emmanuel J. Candès, Yingying Fan, Lucas Janson, and Jinchi Lv. 2016. Panning for Gold: Model-free Knockoffs for High-dimensional Controlled Variable Selection.
- Chenguang Dai, Buyu Lin, Xin Xing, and Jun S Liu. 2020. False Discovery Rate Control via Data Splitting. *arXiv preprint arXiv:2002.08542* (2020).
- Edgar Dobriban and Yue Sheng. 2018. Distributed linear regression by averaging. *arXiv preprint arXiv:1810.00412* (2018).
- Kristen Emery and Uri Keich. 2019. Controlling the FDR in variable selection via multiple knockoffs. *arXiv: Methodology* (2019).
- Ankit Garg, Tengyu Ma, and Huy Nguyen. 2014. On communication cost of distributed statistical estimation and dimensionality. In *Advances in Neural Information Processing Systems*. 2726–2734.
- Atefeh Gilani, Selma Belhadj Amor, Sadaf Salehkalaibar, and Vincent YF Tan. 2019. Distributed hypothesis testing with privacy constraints. *Entropy* 21, 5 (2019), 478.
- Benjamin Hofner, Luigi Boccuto, and Markus Göker. 2015. Controlling false discoveries in high-dimensional situations: boosting with stability selection. In *BMC Bioinformatics*.
- Michael I. Jordan, Jason D. Lee, and Yun Yang. 2019. Communication-Efficient Distributed Statistical Inference. *J. Amer. Statist. Assoc.* 114, 526 (2019), 668–681. <https://doi.org/10.1080/01621459.2018.1429274>
arXiv:<https://doi.org/10.1080/01621459.2018.1429274>
- Alon Kipnis and John C Duchi. 2019. Mean estimation from one-bit measurements. *arXiv preprint arXiv:1901.03403* (2019).
- Nicolai Meinshausen and Peter Bühlmann. 2008. Stability Selection.
- Tuan-Binh Nguyen, Jérôme-Alexis Chevalier, Bertrand Thirion, and Sylvain Arlot. 2020. Aggregation of Multiple Knockoffs. arXiv:math.ST/2002.09269

- A. Ramdas, J. Chen, M. J. Wainwright, and M. I. Jordan. 2017. QuTE: Decentralized multiple testing on sensor networks with false discovery rate control. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. 6415–6421.
- Rajen Dinesh Shah and Richard J. Samworth. 2011. Variable selection with error control: another look at stability selection.
- Sreejith Sreekumar, Deniz Gündüz, and Asaf Cohen. 2018. Distributed hypothesis testing under privacy constraints. In *2018 IEEE Information Theory Workshop (ITW)*. IEEE, 1–5.
- Weijie Su, Junyang Qian, and Linxi Liu. 2015. Communication-Efficient False Discovery Rate Control via Knockoff Aggregation. (2015). *arXiv:stat.ML/1506.05446*
- Fang Xie and Johannes Lederer. 2019. Aggregated False Discovery Rate Control. *arXiv preprint arXiv:1907.03807* (2019).
- Xin Xing, Zhigen Zhao, and Jun S. Liu. 2019. Controlling False Discovery Rate Using Gaussian Mirrors. *Technical Report* (2019).
- Z Zhang and T Berger. 1988. Estimation via compressed information. (1988).
- Tianqi Zhao, Mladen Kolar, and Han Liu. 2014. A general framework for robust testing and confidence regions in high-dimensional quantile regression. *arXiv preprint arXiv:1412.8724* (2014).