Clipper: p-value-free FDR control on high-throughput data from two conditions

Xinzhou Ge[1,†], Yiling Elaine Chen[1,†], Dongyuan Song[2], MeiLu McDermott[3,4], Kyla Woyshner[3], Antigoni Manousopoulou[3], Wei Li[5], Leo D. Wang[3], and Jingyi Jessica Li[1,2,6,7,8,*]

[1]Department of Statistics, University of California, Los Angeles, CA 90095
[2]Interdepartmental Program in Bioinformatics, University of California, Los Angeles, CA 90095
[3]Beckman Research Institute, City of Hope National Medical Center, Duarte, CA 91010
[4]The Quantitative and Computational Biology section, University of Southern California, Los Angeles, CA 90089
[5]Division of Computational Biomedicine, Department of Biological Chemistry, School of Medicine, University of California, Irvine, CA 92697
[6]Department of Human Genetics, University of California, Los Angeles, CA 90095
[7]Department of Computational Medicine, University of California, Los Angeles, CA 90095
[8]Department of Biostatistics, University of California, Los Angeles, CA 90095
[†]These authors contributed equally to this work.
[*]To whom correspondence should be addressed. Email: jli@stat.ucla.edu

Abstract

High-throughput biological data analysis commonly involves the identification of "interesting" features (e.g., genes, genomic regions, and proteins), whose values differ between two conditions, from numerous features measured simultaneously. To ensure the reliability of such analysis, the most widely-used criterion is the false discovery rate (FDR), the expected proportion of uninteresting features among the identified ones. Existing bioinformatics tools primarily control the FDR based on p-values. However, obtaining valid p-values relies on either reasonable assumptions of data distribution or large numbers of replicates under both conditions, two requirements that are often unmet in biological studies. To address this issue, we propose Clipper, a general statistical framework for FDR control without relying on p-values or specific data distributions. Clipper is applicable to identifying both enriched and differential features from high-throughput biological data of diverse types. In comprehensive simulation and real-data benchmarking, Clipper outperforms existing generic FDR control methods and specific bioinformatics tools designed for various tasks, including peak calling from ChIP-seq data, differentially expressed gene identification from RNA-seq data, differentially interacting chromatin region identification from Hi-C data, and peptide identification from mass spectrometry data. Notably, our benchmarking results for peptide identification are based on the first mass spectrometry data standard that has a realistic dynamic range. Our results demonstrate Clipper's flexibility and reliability for FDR control, as well as its broad applications in high-throughput data analysis.

1 Introduction

High-throughput technologies are widely used to provide numerical measurements of system-wide biological features, such as genes, genomic regions, and proteins ("high-throughput" means the number of features is large, at least in thousands). The most common goal of analyzing high-throughput data is to contrast two conditions so as to reliably screen "interesting features", where "interesting" means "enriched" or "differential". "Enriched features" are defined to have higher expected measurements (without measurement errors) under the experimental/treatment condition than the background condition, i.e.,

the negative control. The detection of enriched features is called "enrichment analysis." For example, common enrichment analyses include calling protein-binding sites in a genome from chromatin immunoprecipitation sequencing (ChIP-seq) data [1, 2] and identifying peptides from mass spectrometry (MS) data [3]. In contrast, "differential features" are defined to have different expected measurements between two conditions, and their detection is called "differential analysis." For example, popular differential analyses include the identification of differentially expressed genes (DEGs) from genome-wide gene expression data (e.g., microarray and RNA sequencing (RNA-seq) data [4–10]) and differentially interacting chromatin regions (DIRs) from Hi-C data [11–13]. In most scientific research, the interesting features only constitute a small proportion of all features, and the remaining majority is referred to as "uninteresting features."

The identified features, also called the "discoveries" from differential or enrichment analysis, are subject to further investigation and validation. Hence, to reduce experimental validation that are often laborious and expensive, researchers demand reliable discoveries that contain few false discoveries. Accordingly, the false discovery rate (FDR) [14] was developed as a statistical criterion for ensuring the reliability of discoveries. Technically, the FDR is defined as the expected proportion of uninteresting features among the discoveries under the frequentist paradigm. In parallel, under the Bayesian paradigm, other criteria were developed, including the Bayesian false discovery rate [15], the local false discovery rate [16], and the local false sign rate [17]. Among all these frequentist and Bayesian criteria, the FDR is the dominant criterion for setting thresholds in biological data analysis [1, 10, 18–24] and is thus the focus of this paper.

FDR control refers to the goal of finding discoveries such that the FDR is under a pre-specified threshold (e.g., 0.05). Existing computational methods for FDR control primarily rely on p-values, one per feature. Among the p-value-based methods, the most classic and popular ones are the Benjamini-Hochberg (BH) procedure [14] and the Storey's q-value [25]; later development introduced methods that incorporate feature weights [26] or covariates (e.g., independent hypothesis weighting (IHW) [27], adaptive p-value thresholding [28], and Boca and Leek's FDR regression [29]) to boost the detection power. All these methods set a p-value cutoff based on the pre-specified FDR threshold. However, the calculation of p-values requires either distributional assumptions or large numbers of replicates, which are often questionable or unachievable in biological data analysis. We categorize existing ways of calculating p-values into "paired" and "pooled" approaches, which are most commonly used in bioinformatics tools. Below we describe how each approach is vulnerable to unmet distributional assumptions or small numbers of replicates.

The paired approach, which has been widely used to detect DEGs and protein-binding sites [1, 2, 4, 5], examines one feature at a time and compares the feature's measurements between two conditions ("paired" means that each feature has two sets of measurements) using a statistical test. However, a valid test always requires a not-too-small number of replicates. The reason is that, without enough replicates, parametric tests cannot have their specific distributions verified, and non-parametric tests become powerless. While biomedical researchers understand that non-parametric tests are unsuitable when the number of replicates is small, they often neglect the distribution misspecification issue with parametric tests. An example is the detection of protein-binding sites from ChIP-seq data. A common assumption is that ChIP-seq read counts in a genomic region (i.e., a feature) follow the Poisson distribution [1, 2], which implies that the counts have a variance equal to their mean. However, if only two replicates are available, it is impossible to check whether this Poisson distribution is reasonably specified. Further, another issue with the paired approach—regardless of the test being parametric or non-parametric—is the misspecification of a two-sample test as a one-sample test, which ignores the sampling randomness of replicates under one condition. This issue is implicit but widespread in bioinformatics methods [1, 30].

The pooled approach requires that the two conditions are specified as the background and experimental/treatment conditions, respectively, and that the two conditions have the same number of replicates. It is often used when both conditions only have one replicate, and its important application in proteomics is to identify peptide sequences from MS data [31]. Technically, for every feature, it defines the test statistic as the feature's average measurement under the experimental condition. To form a null distribution for the test statistic, it pools all features' average measurements under the background condition, and that is why we call it "the pooled approach." Finally, it calculates a p-value for every feature based on the feature's observed test statistic and the null distribution. Although it does not have apparent assumptions, its validity requires an implicit assumption that all features come from the same distribution under the background condition. However, even though this assumption holds in the aforementioned proteomics application [32], it is often violated due to the existence of feature heterogeneity in many other biological applications; for example, in the identification of protein-binding sites, different genomic regions exhibit diverse non-specific binding affinities, [33], so these regions' ChIP-seq read counts do not follow the same distribution under the background condition. When such violation occurs, the pooled approach would output invalid p-values and leads to unreliable FDR control. In addition, the pooled approach requires a sufficiently large number of features (at least in thousands) to estimate the null distribution well; otherwise, the p-values would not be well-calibrated.

Due to the limitations of both paired and pooled approaches in high-throughput biological data analysis, bioinformatics tools often output ill-posed p-values. This is evidenced in serious concerns in the scientific community about the widespread miscalculation and misuses of p-values [34]. As a result, bioinformatics tools relying on questionable p-values either cannot reliably control the FDR to a target level [23] or lack power to make discoveries. Therefore, p-value-free control of FDR is desirable, as it would make data analysis more transparent and thus improve the reproducibility of scientific research.

Here we propose a general statistical framework Clipper to provide reliable FDR control for high-throughput biological data analysis, without using p-values or relying on specific data distributions. Clipper defines a contrast score, as replacement of a p-value, to summarize and contrast each feature's measurements between two conditions. Then, as its name suggests, Clipper establishes a cutoff on features' contrast scores to identify a small proportion of interesting features. Clipper is a flexible framework that only requires a minimal input: all features' measurements under two conditions and a target FDR threshold (e.g., 0.05), and Clipper is applicable to both enrichment and differential analyses (Fig. 1). To verify Clipper's performance, we designed comprehensive simulation studies to benchmark Clipper against existing generic FDR control methods. We also benchmarked Clipper against bioinformatics tools in studies including peak calling from ChIP-seq data, peptide identification from mass spectrometry data, DEG identification from RNA-seq data, and DIR identification from Hi-C data. Notably, our benchmarking results for peptide identification are based on the first MS data standard that has a realistic dynamic range. In summary, Clipper is a robust and flexible framework that works for high-throughput data with various characteristics, including measurement distributions, replicate numbers (from one to multiple), and outlier existence.

2 The Clipper methodology

Clipper is a flexible framework that reliably controls the FDR without using p-values in high-throughput data analysis with two conditions. Clipper has two functionalities: (I) enrichment analysis, which identifies the "interesting" features that have higher expected measurements (i.e., true signals) under the experimental condition than the background, a.k.a. negative control condition (if the goal is to identify the interesting features with smaller expected measurements under the experimental condition, enrich-

ment analysis can be applied after the values are negated); (II) differential analysis, which identifies the interesting features that have different expected measurements between the two conditions. For both functionalities, uninteresting features are defined as those that have equal expected measurements under the two conditions.

Clipper only relies on two fundamental statistical assumptions of biological data analysis: (1) measurement errors (i.e., differences between measurements and their expectations, which include biological signals and batch effects) are independent across all features and experiments; (2) every uninteresting feature has measurement errors identically distributed across all experiments. These two assumptions are used in almost all bioinformatics tools and commonly referred to as the "measurement model" in statistical genomics [35].

Clipper consists of two major steps: construction and thresholding of contrast scores. First, Clipper computes contrast scores, one per feature, as summary statistics that reflect the extent to which features are interesting. Second, Clipper establishes a contrast-score cutoff and calls as discoveries the features whose contrast scores exceed the cutoff.

In the following subsections, we will first introduce notations and assumptions used in Clipper. Then we will detail how Clipper works and discuss its theoretical guarantee in three analysis tasks: the enrichment analysis with equal numbers of replicates under two conditions, the enrichment analysis with different numbers of replicates under two conditions, and the differential analysis.

2.1 Notations and assumptions

To facilitate our discussion, we first introduce the following mathematical notations. For two random vectors $\boldsymbol{X} = (X_1, \ldots, X_m)^\top$ and $\boldsymbol{Y} = (Y_1, \ldots, Y_n)^\top$, or two sets of random variables $\mathcal{X} = \{X_1, \ldots, X_m\}$ and $\mathcal{Y} = \{Y_1, \ldots, Y_n\}$, we write $\boldsymbol{X} \perp \boldsymbol{Y}$ or $\mathcal{X} \perp \mathcal{Y}$ if X_i is independent of Y_j for all $i = 1, \ldots, m$ and $j = 1, \ldots, n$. To avoid confusion, we use $\mathrm{card}(A)$ to denote the cardinality of a set A and $|c|$ to denote the absolute value of a scalar c. We define $a \vee b = \max(a, b)$.

While the differential analysis treats the two conditions symmetric, the enrichment analysis requires one condition to be the experimental condition (i.e., the condition of interest) and the other to be the background condition (i.e., the negative control). For simplicity, we use the same set of notations for both analyses.

Clipper only requires two inputs: the target FDR threshold $q \in (0, 1)$ and the input data. Regarding the input data, we use d to denote the number of features with measurements under two conditions, and we use m and n to denote the numbers of replicates under the two conditions. For each feature $j = 1, \ldots, d$, we use $\boldsymbol{X}_j = (X_{j1}, \ldots, X_{jm})^\top \in \mathbb{R}^m_{\geq 0}$ and $\boldsymbol{Y}_j = (Y_{j1}, \ldots, Y_{jn})^\top \in \mathbb{R}^n_{\geq 0}$ to denote its measurements under the two conditions, where $\mathbb{R}_{\geq 0}$ denotes the set of non-negative real numbers. We assume that all measurements are non-negative, as in the case of most high-throughput experiments. (If this assumption does not hold, transformations can be applied to make data satisfy this assumption.)

Clipper has the following assumptions on the joint distribution of $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_d, \boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_d$. For $j = 1, \ldots, d$, Clipper assumes that $X_{j1}, \ldots, X_{jm}$ are identically distributed, so are $Y_{j1}, \ldots, Y_{jn}$. Let $\mu_{Xj} = \mathbb{E}[X_{j1}]$ and $\mu_{Yj} = \mathbb{E}[Y_{j1}]$ denote the expected measurement of feature j under the two condition, respectively. Then conditioning on $\{\mu_{X_j}\}_{j=1}^d$ and $\{\mu_{Y_j}\}_{j=1}^d$,

$$\begin{aligned} &X_{j1}, \cdots, X_{jm}, Y_{j1}, \cdots, Y_{jn} \text{ are mutually independent}\,; \\ &\boldsymbol{X}_j \perp \boldsymbol{X}_k, \boldsymbol{Y}_j \perp \boldsymbol{Y}_k \text{ and } \boldsymbol{X}_j \perp \boldsymbol{Y}_k\,,\ \forall j, k = 1, \ldots, d. \end{aligned} \tag{1}$$

An enrichment analysis aims to identify interesting features with $\mu_{Xj} > \mu_{Yj}$ (with $\boldsymbol{X}_j$ and $\boldsymbol{Y}_j$ defined as the measurements under the experimental and background conditions, respectively), while a differential

analysis aims to call interesting features with $\mu_{Xj} \neq \mu_{Yj}$. We define $\mathcal{N} := \{j : \mu_{Xj} = \mu_{Yj}\}$ as the set of uninteresting features and denote $N := \mathrm{card}(\mathcal{N})$. In both analyses, Clipper further assumes that an uninteresting feature j satisfies

$$X_{j1}, \cdots, X_{jm}, Y_{j1}, \cdots, Y_{jn} \text{ are identically distributed}, \forall j \in \mathcal{N}. \tag{2}$$

To construct contrast scores, Clipper uses two summary statistics $t(\cdot,\cdot) : \mathbb{R}^m_{\geq 0} \times \mathbb{R}^n_{\geq 0} \to \mathbb{R}$ to extract data information regarding whether a feature is interesting or not:

$$t^{\mathrm{diff}}(\boldsymbol{x}, \boldsymbol{y}) := \bar{x} - \bar{y}\,; \tag{3}$$

$$t^{\max}(\boldsymbol{x}, \boldsymbol{y}) := \max(\bar{x}, \bar{y}) \cdot \mathrm{sign}(\bar{x} - \bar{y})\,, \tag{4}$$

where $\boldsymbol{x} = (x_1, \ldots, x_m)^\top \in \mathbb{R}^m_{\geq 0}$, $\boldsymbol{y} = (y_1, \ldots, y_n)^\top \in \mathbb{R}^n_{\geq 0}$, $\bar{x} = \sum_{i=1}^m x_i/m$, $\bar{y} = \sum_{i=1}^n y_i/n$, and $\mathrm{sign}(\cdot) : \mathbb{R} \to \{-1, 0, 1\}$ with $\mathrm{sign}(x) = 1$ if $x > 0$, -1 if $x < 0$, and 0 otherwise.

2.2 Enrichment analysis with equal numbers of replicates ($m = n$)

Under the enrichment analysis, we assume that $\boldsymbol{X}_j \in \mathbb{R}^m_{\geq 0}$ and $\boldsymbol{Y}_j \in \mathbb{R}^n_{\geq 0}$ are the measurements of feature j, $j = 1, \ldots, d$, under the experimental and background conditions with m and n replicates, respectively. We start with the simple case when $m = n$. Clipper defines a contrast score C_j of feature j in one of two ways:

$$C_j := t^{\mathrm{diff}}(\boldsymbol{X}_j, \boldsymbol{Y}_j) \qquad \textbf{difference contrast score}\,, \tag{5}$$

or

$$C_j := t^{\max}(\boldsymbol{X}_j, \boldsymbol{Y}_j) \qquad \textbf{maximum contrast score}\,. \tag{6}$$

Accordingly, a large positive value of C_j bears evidence that $\mu_{Xj} > \mu_{Yj}$. Motivated by Barber and Candès [36], Clipper proposes the following procedure to control the FDR under the target level $q \in (0, 1)$.

Definition 1 (Barber-Candès (BC) procedure for thresholding contrast scores [36]) *Given contrast scores $\{C_j\}_{j=1}^d$, $\mathcal{C} = \{|C_j| : C_j \neq 0\,;\ j = 1, \ldots, d\}$ is defined as the set of non-zero absolute values of C_j's. The BC procedure finds a contrast-score cutoff T^{BC} based on the target FDR threshold $q \in (0, 1)$ as*

$$T^{\mathrm{BC}} := \min\left\{t \in \mathcal{C} : \frac{\mathrm{card}(\{j : C_j \leq -t\}) + 1}{\mathrm{card}(\{j : C_j \geq t\}) \vee 1} \leq q\right\} \tag{7}$$

and outputs $\{j : C_j \geq T^{\mathrm{BC}}\}$ as discoveries.

Theorem 1 *Suppose that the input data satisfy the Clipper assumptions (1)–(2) and $m = n$. Then for any $q \in (0, 1)$ and either definition of constrast scores in (5) or (6), the contrast-score cutoff T^{BC} found by the BC procedure guarantees that the discoveries have the FDR under q:*

$$\mathrm{FDR} = \mathbb{E}\left[\frac{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\}\right)}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{BC}}\}) \vee 1}\right] \leq q\,,$$

where $\mathcal{N} = \{j : \mu_{Xj} = \mu_{Yj}\}$ denotes the set of uninteresting features.

The proof of Theorem 1 (Supplementary Section S6) requires two key ingredients: Lemma 1, which states important properties of contrast scores, and Lemma 2 from [37], which states a property of a

Bernoulli process with independent but not necessarily identically distributed random variables. The cutoff T^{BC} can be viewed as a stopping time of a Bernoulli process.

Lemma 1 *Suppose that the input data that satisfy the Clipper assumption (1)–(2) and $m = n$, and that Clipper constructs contrast scores $\{C_j\}_{j=1}^d$ based on (5) or (6). Denote $S_j = \mathrm{sign}(C_j) \in \{-1, 0, 1\}$. Then $\{S_j\}_{j=1}^d$ satisfy the following properties:*

(a) $S_1, \ldots, S_d$ are mutually independent ;

(b) $\mathbb{P}(S_j = 1) = \mathbb{P}(S_j = -1)$ for all $j \in \mathcal{N}$;

(c) $\{S_j\}_{j \in \mathcal{N}} \perp \mathcal{C}$.

Lemma 2 *Suppose that $Z_1, \ldots, Z_d$ are independent with $Z_j \sim \mathrm{Bernoulli}(\rho_j)$, and $\min_j \rho_j \geq \rho > 0$. Let J be a stopping time in reverse time with respect to the filtration $\{\mathcal{F}_j\}$, where*

$$\mathcal{F}_j = \sigma\left(\{(Z_1 + \cdots + Z_j),\, Z_{j+1},\, \cdots,\, Z_d\}\right), \tag{8}$$

with $\sigma(\cdot)$ denoting a σ-algebra. Then

$$\mathbb{E}\left[\frac{1+J}{1+Z_1+\cdots+Z_J}\right] \leq \rho^{-1}.$$

Here we give a brief intuition about how Lemma 2 bridges Lemma 1 and Theorem 1 for FDR control. First we note that the false discovery proportion (FDP), whose expectation is the FDR, satisfies

$$\begin{aligned}
\mathrm{FDP} &:= \frac{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\}\right)}{\mathrm{card}\left(\{j : C_j \geq T^{\mathrm{BC}}\}\right) \vee 1} \\
&= \frac{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\}\right)}{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}\right) + 1} \cdot \frac{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}\right) + 1}{\mathrm{card}\left(\{j : C_j \geq T^{\mathrm{BC}}\}\right) \vee 1} \\
&\leq \frac{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\}\right)}{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}\right) + 1} \cdot \frac{\mathrm{card}\left(\{j : C_j \leq -T^{\mathrm{BC}}\}\right) + 1}{\mathrm{card}\left(\{j : C_j \geq T^{\mathrm{BC}}\}\right) \vee 1} \\
&\leq \frac{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\}\right)}{\mathrm{card}\left(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}\right) + 1} \cdot q\,,
\end{aligned}$$

where the last inequality follows from the definition of T^{BC} (7).

By its definition, if T^{BC} exists, it is positive. This implies that Clipper would never call the features with $C_j = 0$ as discoveries. Here we sketch the idea of proving Theorem 1 by considering a simplified case where $\mathcal{C}$ is fixed instead of being random; that is, we assume the features with non-zero contrast scores to be known. Then, without loss of generality, we assume $\mathcal{C} = \{1, \ldots, d\}$. Then we order the absolute values of uninteresting features' contrast scores, i.e., elements in $\{|C_j| : j \in \mathcal{N}\}$, from the largest to the smallest, denoted by $|C_{(1)}| \geq |C_{(2)}| \geq \cdots \geq |C_{(N)}|$. Let $J = \sum_{j \in \mathcal{N}} \mathbb{1}\left(|C_j| \geq T^{\mathrm{BC}}\right)$, the number of uninteresting features whose contrast scores have absolute values no less than T^{BC}. When $J > 0$, $|C_{(1)}| \geq \cdots \geq |C_{(J)}| \geq T^{\mathrm{BC}}$. Define $Z_k = \mathbb{1}\left(C_{(k)} < 0\right)$, $k = 1, \ldots, N$. Then for each order k, the following holds

$$\begin{aligned}
C_{(k)} \geq T^{\mathrm{BC}} &\Longleftrightarrow \left|C_{(k)}\right| \geq T^{\mathrm{BC}} \text{ and } C_{(k)} > 0 \Longleftrightarrow k \leq J \text{ and } Z_k = 0\,; \\
C_{(k)} \leq -T^{\mathrm{BC}} &\Longleftrightarrow \left|C_{(k)}\right| \geq T^{\mathrm{BC}} \text{ and } C_{(k)} < 0 \Longleftrightarrow k \leq J \text{ and } Z_k = 1\,.
\end{aligned}$$

Then the upper bound of FDP becomes

$$\begin{aligned}\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}) + 1} \cdot q &= \frac{\sum_{k=1}^{N} \mathbb{1}\left(C_{(k)} \geq T^{\mathrm{BC}}\right)}{1 + \sum_{k=1}^{N} \mathbb{1}\left(C_{(k)} \leq -T^{\mathrm{BC}}\right)} \cdot q \\ &= \frac{\sum_{k=1}^{J} \mathbb{1}\left(C_{(k)} \geq T^{\mathrm{BC}}\right)}{1 + \sum_{k=1}^{J} \mathbb{1}\left(C_{(k)} \leq -T^{\mathrm{BC}}\right)} \cdot q \\ &= \frac{(1 - Z_1) + \cdots + (1 - Z_J)}{1 + Z_1 + \cdots + Z_J} \cdot q \\ &= \left(\frac{1 + J}{1 + Z_1 + \cdots + Z_J} - 1\right) \cdot q\,.\end{aligned}$$

By Lemma 1(a)–(b), $Z_k \overset{\text{i.i.d.}}{\sim} \mathrm{Bernoulli}(0.5)$, which together with Lemma 1(c) satisfy the condition of Lemma 2 and make $\rho = 0.5$. Then by Lemma 2, we have

$$\mathrm{FDR} = \mathbb{E}[\mathrm{FDP}] \leq \mathbb{E}\left[\frac{1 + J}{1 + Z_1 + \cdots + Z_J} - 1\right] \cdot q \leq (\rho^{-1} - 1) \cdot q = q\,,$$

which is the statement of Theorem 1. The complete proof of Theorem 1 is in Supplementary Section S6.

2.2.1 An optional, heuristic fix if the BC procedure makes no discoveries

Although the BC procedure has theoretical guarantee of FDR control, it lacks power when the number of replicates $m = n$, the target FDR threshold q, and the number of features d are all small (e.g., $m = n = 1$, $q = 0.01$ and $d = 1000$ in Fig. 2). As a result, the BC procedure may lead to no discoveries. In that case, Clipper implements a heuristic fix—an approximate p-value Benjamini-Hochberg (aBH) procedure—to increase the power. The aBH procedure constructs an empirical null distribution of contrast scores by additionally assuming that uninteresting features' contrast scores follow a symmetric distribution around zero; it then computes approximate p-values of features based on the empirical null distribution, and finally it uses the BH procedure [14] to threshold the approximate p-values.

Definition 2 (The aBH procedure) *Given contrast scores $\{C_j\}_{j=1}^{d}$, an empirical null distribution is defined on $\mathcal{E} := \{C_j : C_j < 0; j = 1, \ldots, d\} \cup \{-C_j : C_j < 0; j = 1, \ldots, d\}$. The aBH procedure defines the approximate p-value of feature j as*

$$p_j := \frac{\sum_{c \in \mathcal{E}} \mathbb{1}(c \geq C_j)}{\mathrm{card}(\mathcal{E}) \vee 1}\,.$$

Then it applies the BH procedure with the target FDR threshold q to $\{p_j\}_{j=1}^{d}$ to call discoveries.

2.3 Enrichment analysis with any numbers of replicates m and n

When $m \neq n$, the BC procedure cannot guarantee FDR control because Lemma 1 no longer holds. To control the FDR in a more general setting ($m = n$ or $m \neq n$), Clipper constructs contrast scores via permutation of replicates across conditions. The idea is that, after permutation, every feature becomes uninteresting and can serve as its own negative control.

Definition 3 (Permutation) *We define σ as permutation, i.e., a bijection from the set $\{1, \cdots, m + n\}$*

onto itself, and we rewrite the data $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_d, \boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_d$ into a matrix $\mathbf{W}$:

$$\mathbf{W} = \begin{bmatrix} W_{11} & \cdots & W_{1m} & W_{1(m+1)} & \cdots & W_{1(m+n)} \\ & \vdots & & & \vdots & \\ W_{d1} & \cdots & W_{dm} & W_{d(m+1)} & \cdots & W_{d(m+n)} \end{bmatrix} := \begin{bmatrix} X_{11} & \cdots & X_{1m} & Y_{11} & \cdots & Y_{1n} \\ & \vdots & & & \vdots & \\ X_{d1} & \cdots & X_{dm} & Y_{d1} & \cdots & Y_{dn} \end{bmatrix}.$$

We then apply σ to permute the columns of $\mathbf{W}$ and obtain

$$\mathbf{W}_\sigma := \begin{bmatrix} W_{1\sigma(1)} & \cdots & W_{1\sigma(m)} & W_{1\sigma(m+1)} & \cdots & X_{1\sigma(m+n)} \\ & \vdots & & & \vdots & \\ W_{d\sigma(1)} & \cdots & W_{d\sigma(m)} & W_{d\sigma(m+1)} & \cdots & W_{d\sigma(m+n)} \end{bmatrix},$$

from which we obtain the permuted measurements $\left\{(\boldsymbol{X}_j^\sigma, \boldsymbol{Y}_j^\sigma)\right\}_{j=1}^d$, where

$$\begin{aligned} \boldsymbol{X}_j^\sigma &:= \left(W_{j\sigma(1)}, \ldots, W_{j\sigma(m)}\right)^\top, \\ \boldsymbol{Y}_j^\sigma &:= \left(W_{j\sigma(m+1)}, \ldots, W_{j\sigma(m+n)}\right)^\top. \end{aligned} \tag{9}$$

In the enrichment analysis, if two permutations σ and σ' satisfy that

$$\{\sigma(1), \cdots, \sigma(m)\} = \{\sigma'(1), \cdots, \sigma'(m)\},$$

then we define σ and σ' to be in one equivalence class. That is, permutations in the same equivalence class lead to the same division of $m+n$ replicates (from the two conditions) into two groups with sizes m and n. In total, there are $\binom{m+n}{m}$ equivalence classes of permutations.

We define σ_0 as the identity permutation such that $\sigma_0(i) = i$ for all $i \in \{1, \cdots, m+n\}$. In addition, Clipper randomly samples h equivalence classes $\sigma_1, \ldots, \sigma_h$ with equal probabilities without replacement from the other $h_{\max} := \binom{m+n}{m} - 1$ equivalence classes (after excluding the equivalence class containing σ_0). Note that $h_{\max}$ is the maximum value h can take.

Clipper then obtains $\left\{(\boldsymbol{X}_j^{\sigma_0}, \boldsymbol{Y}_j^{\sigma_0}), (\boldsymbol{X}_j^{\sigma_1}, \boldsymbol{Y}_j^{\sigma_1}), \cdots, (\boldsymbol{X}_j^{\sigma_h}, \boldsymbol{Y}_j^{\sigma_h})\right\}_{j=1}^d$, where $(\boldsymbol{X}_j^{\sigma_\ell}, \boldsymbol{Y}_j^{\sigma_\ell})$ are the permuted measurements based on σ_ℓ, $\ell = 0, \ldots, h$. Then Clipper computes $T_j^{\sigma_\ell} := t^{\text{diff}}(\boldsymbol{X}_j^{\sigma_\ell}, \boldsymbol{Y}_j^{\sigma_\ell})$ to indicate the degree of "interestingness" of feature j reflected by $(\boldsymbol{X}_j^{\sigma_\ell}, \boldsymbol{Y}_j^{\sigma_\ell})$. Note that Clipper chooses t^{diff} instead of $t^{\max}$ because empirical evidence shows that t^{diff} leads to better power. Sorting $\{T_j^{\sigma_\ell}\}_{\ell=0}^h$ gives

$$T_j^{(0)} \geq T_j^{(1)} \geq \cdots \geq T_j^{(h)}.$$

Then Clipper defines the contrast score of feature j, $j = 1, \ldots, d$, in one of two ways:

$$C_j := \begin{cases} T_j^{(0)} - T_j^{(1)} & \text{if } T_j^{(0)} = T_j^{\sigma_0} \\ T_j^{(1)} - T_j^{(0)} & \text{otherwise} \end{cases} \qquad \textbf{difference contrast score}, \tag{10}$$

or

$$C_j := \begin{cases} \left|T_j^{(0)}\right| & \text{if } T_j^{(0)} = T_j^{\sigma_0} > T_j^{(1)} \\ 0 & \text{if } T_j^{(0)} = T_j^{(1)} \\ -\left|T_j^{(0)}\right| & \text{otherwise} \end{cases} \qquad \textbf{maximum contrast score}. \tag{11}$$

The intuition behind the contrast scores is that, if $C_j < 0$, then $\mathbb{1}(T_j^{(0)} = T_j^{\sigma_0}) = 0$, which means that at least one of $T_j^{\sigma_1}, \ldots, T_j^{\sigma_h}$ (after random permutation) is greater than $T_j^{\sigma_0}$ calculated from the original data (identity permutation), suggesting that feature j is likely an uninteresting feature in enrichment

analysis. Motivated by Gimenez and Zou [38], we propose the following procedure for Clipper to control the FDR under the target level $q \in (0, 1)$.

Definition 4 (Gimenez-Zou (GZ) procedure for thresholding contrast scores [38]) *Given $h \in \{1, \cdots, h_{\max}\}$ and contrast scores $\{C_j\}_{j=1}^d$, $\mathcal{C} = \{|C_j| : C_j \neq 0 \,;\, j = 1, \ldots, d\}$ is defined as the set of non-zero absolute values of C_j's. The GZ procedure finds a contrast-score cutoff T^{GZ} based on the target FDR threshold $q \in (0, 1)$ as:*

$$T^{\mathrm{GZ}} := \min \left\{ t \in \mathcal{C} : \frac{\frac{1}{h} + \frac{1}{h}\mathrm{card}\left(\{j : C_j \leq -t\}\right)}{\mathrm{card}\left(\{j : C_j \geq t\}\right) \vee 1} \leq q \right\} \tag{12}$$

and outputs $\left\{j : C_j \geq T^{\mathrm{GZ}}\right\}$ as discoveries.

Theorem 2 *Suppose that the input data that satisfy the Clipper assumption (1)–(2). Then for any $q \in (0, 1)$ and either definition of contrast scores in (10) or (11), the contrast-score cutoff T^{GZ} found by the GZ procedure (12) guarantees that the discoveries have the FDR under q:*

$$\mathrm{FDR} = \mathbb{E}\left[\frac{\mathrm{card}\left(\left\{j \in \mathcal{N} : C_j \geq T^{\mathrm{GZ}}\right\}\right)}{\mathrm{card}\left(\{j : C_j \geq T^{\mathrm{GZ}}\}\right) \vee 1}\right] \leq q\,,$$

where $\mathcal{N}$ denotes the set of uninteresting features.

The proof of Theorem 2 (Supplementary Section S6) is similar to that of Theorem 1 and requires two key ingredients: Lemma 2, which is also used in the proof of Theorem 1, and Lemma 3, which is similar to Lemma 1 and is about the properties of signs of $\{C_j\}_{j=1}^d$. The cutoff T^{GZ} can also be viewed as a stopping time of a Bernoulli process.

Lemma 3 *For input data that satisfy the Clipper assumption (1) and (2), Clipper constructs contrast scores $\{C_j\}_{j=1}^d$ based on (11) or (10). Denote $S_j = \mathrm{sign}\left(C_j\right) \in \{-1, 0, 1\}$. Then $\{S_j\}_{j=1}^d$ and $\{C_j\}_{j=1}^d$ satisfy the following properties:*

(a) $S_1, \ldots, S_d$ are mutually independent ;

(b) $\mathbb{P}(S_j = 1) \leq \frac{1}{h+1}$ for all $j \in \mathcal{N}$;

(c) $\{S_j\}_{j \in \mathcal{N}} \perp \mathcal{C}$.

We note that the GZ procedure is also applicable to the enrichment analysis with equal numbers of replicates, i.e., $m = n$ (Section 2.2). We will compare the GZ procedure against the BC procedure in our results.

2.4 Differential analysis with $m + n > 2$

For differential analysis, Clipper also uses permutation to construct contrast scores. When $m \neq n$, the equivalence classes of permutations are defined the same as for the enrichment analysis with $m \neq n$ (Section 2.3). When $m = n$, there is a slight change in the definition of equivalence classes of permutations: if σ and σ' satisfy that

$$\{\sigma(1), \cdots, \sigma(m)\} = \{\sigma'(1), \cdots, \sigma'(m)\} \text{ or } \{\sigma'(m+1), \cdots, \sigma'(2m)\}\,,$$

then we say that σ and σ' are in one equivalence class. In total, there are $\binom{2m}{m}/2$ equivalence classes of permutations. Hence, to have more than one equivalence class, we cannot perform differential analysis with $m = n = 1$; in other words, the total number of replicates $m + n$ must be at least 3.

Then Clipper randomly samples $\sigma_1, \ldots, \sigma_h$ with equal probabilities without replacement from the $h_{\max} := \binom{2m}{m}/2 - 1$ equivalence classes that exclude the class containing σ_0, i.e., the identity permutation. Note that $h_{\max}$ is the maximum value h can take. Next, Clipper computes $T_j^{\sigma_\ell} := \left|t^{\text{diff}}(\boldsymbol{X}_j^{\sigma_\ell}, \boldsymbol{Y}_j^{\sigma_\ell})\right|$, where $\boldsymbol{X}_j^{\sigma_\ell}$ and $\boldsymbol{Y}_j^{\sigma_\ell}$ are the permuted data defined in (9), and it defines a C_j as the contrast score of feature j, $j = 1, \ldots, d$ in the same ways as in (10) or (11).

Same as in the enrichment analysis with $m \neq n$, Clipper also uses the GZ procedure [38] to set a cutoff on contrast scores to control the FDR under the target level $q \in (0, 1)$, following Theorem 2.

2.5 Clipper variant algorithms

For nomenclature, we assign the following names to Clipper variant algorithms, each of which combines a contrast score definition with a thresholding procedure.

- **Clipper-diff-BC**: difference contrast score $C_j = t^{\text{diff}}(\boldsymbol{X}_j, \boldsymbol{Y}_j)$ (5) and BC procedure (Definition 1);
- **Clipper-diff-aBH**: difference contrast score $C_j = t^{\text{diff}}(\boldsymbol{X}_j, \boldsymbol{Y}_j)$ and aBH procedure (Definition 2);
- **Clipper-diff-GZ**: difference contrast score $\tau_j = T_j^{(0)} - T_j^{(1)}$ (10) and GZ procedure (Definition 4);
- **Clipper-max-BC**: maximum contrast score $C_j = t^{\max}(\boldsymbol{X}_j, \boldsymbol{Y}_j)$ (6) and BC procedure;
- **Clipper-max-aBH**: maximum contrast score $C_j = t^{\max}(\boldsymbol{X}_j, \boldsymbol{Y}_j)$ and aBH procedure;
- **Clipper-max-GZ**: maximum contrast score $\tau_j = T_j^{(0)}$ (11) and GZ procedure.

2.6 `R` package "Clipper"

In the `R` package `Clipper`, the default implementation is as follows. Based on the power comparison results in Figs. 2, 4, 6, and 8, Clipper uses Clipper-diff-BC as the default algorithm for the enrichment analysis with equal numbers of replicates; when there are no discoveries, Clipper suggests users to increase the target FDR threshold q or to use the Clipper-diff-aBH algorithm with the current q. For the enrichment analysis with different numbers of replicates under two conditions or the differential analysis, Clipper uses the Clipper-max-GZ algorithm by default.

3 Simulation analysis

We designed comprehensive simulation studies to validate the FDR control by Clipper and to benchmark Clipper against existing FDR control methods. Simulation is essential because we can generate numerous datasets from the same distribution with known truths to estimate the FDR, which is not observable from real data. Our simulation covers both enrichment and differential analyses. In enrichment analysis, we consider three "experimental designs": 1vs1 design (one replicate per condition), 2vs1 design (two and one replicates under the experimental and background conditions, respectively), and 3vs3 design (three replicates per condition). In differential analysis, since Clipper requires that at least one condition has two replicates, we only consider the 2vs1 and 3vs3 designs. For each analysis and design, we simulated data from three "distributional families"—Gaussian, Poisson, and negative binomial—for individual features under two "background scenarios" (scenarios of the background condition): homogeneous and heterogeneous. Under the homogeneous scenario, all features' measurements follow the same distribution under the background condition; otherwise, we are under the heterogeneous scenario, which is ubiquitous in applications, e.g., identifying DEGs from RNA-seq data and calling protein-binding sites from ChIP-seq data.

For both enrichment and differential analyses and each simulation setting—a combination of an experimental design, a distributional family, and a background scenario—we first compared Clipper variant algorithms to justify the default Clipper implementation; then we compared Clipper against generic FDR control methods including p-value-based and local-fdr-based methods. Briefly, p-value-based methods use either the paired or pooled approach to calculate p-values of features and then threshold the p-values using the BH procedure (Definition 5) or Storey's qvalue procedure (Definition 6) to make discoveries (Supplementary Section S1.1). Combining a p-value calculation approach and a p-value thresholding procedure, we obtain four p-value-based methods BH-pair, BH-pool, qvalue-pair, and qvalue-pool. In contrast, local-fdr-based methods estimate local false discovery rates (local fdrs) of features and then threshold the local fdrs using the locfdr procedure (Definition 8) to make discoveries (Supplementary Section S1.2). Even though these local-fdr-based methods are not guaranteed to control the FDR, their discoveries are expected to approximately satisfy the FDR control [39, 40]. The estimation of local fdrs takes one of two approaches: empirical nulls, which are estimated parametrically from the test statistic's values likely drawn from the null distribution, and swapping. The resulting two local-fdr-based-methods are referred to as locfdr-emp and locfdr-swap. Supplementary Section S1 provides a detailed explanation of these generic methods and how we implemented them in this work.

Specific to the p-value-based methods, for the paired approach, besides the ideal implementation that uses the correct model to calculate p-values (BH-pair-correct and qvalue-pair-correct), we also consider common implementations that misspecify the distribution (BH-pair-mis and qvalue-pair-mis) or mis-formulate the two-sample test as a one-sample test (BH-pair-2as1 and qvalue-pair-2as1).

To summarize, we compared Clipper against the following implementations of generic FDR control methods:

- **BH-pool** or **qvalue-pool**: p-values calculated by the pooled approach and thresholded by the BH or qvalue procedure
- **BH-pair-correct** or **qvalue-pair-correct**: p-values calculated by the paired approach with the correct model specification and thresholded by the BH or qvalue procedure
- **BH-pair-mis** or **qvalue-pair-mis**: p-values calculated by the paired approach with a misspecified model and thresholded by the BH or qvalue procedure
- **BH-pair-2as1** or **qvalue-pair-2as1**: p-values calculated by the paired approach that mis-formulates a two-sample test as a one-sample test (2as1) and thresholded by the BH or qvalue procedure
- **locfdr-emp**: local fdrs calculated by the empirical null approach and thresholded by the locfdr procedure
- **locfdr-swap**: local fdrs calculated by the swapping approach and thresholded by the locfdr procedure

We demonstrate that, compared with the above ten methods, Clipper controls the FDR and achieves the best power under most settings in both enrichment and differential analyses.

3.1 Simulation settings

In enrichment analysis, we consider three experimental designs.

- For the 1vs1 design where the experimental and background conditions each have one replicate, we simulated data with $d = 1000$ or 10000 features under each distributional family (Gaussian, Poisson, or negative binomial) and each background scenario (homogeneous or heterogeneous)—a total of 12 settings.

- For the 2vs1 design where the experimental and background conditions have two and one replicates respectively, we simulated data under each distributional family (Gaussian, Poisson, or negative binomial) and each background scenario (homogeneous or heterogeneous)—a total of 6 settings.
- For the 3vs3 design where both conditions have three replicates, we simulated data with or without outliers under each distributional family (Gaussian, Poisson, or negative binomial) and each background scenario (homogeneous or heterogeneous)—a total of 12 settings.

In differential analysis, we only consider the 2vs1 and 3vs3 experimental designs, which are the same as in enrichment analysis, because differential analysis requires at least one condition to have more than one replicates.

3.2 Comparison of Clipper variant algorithms

We compared Clipper variant algorithms applicable to each experimental design. Based on the comparison results, we selected a variant algorithm as the default Clipper implementation for each design.

- **1vs1 enrichment analysis.** Under each of the 12 settings, we compared Clipper-diff-BC, Clipper-diff-aBH, Clipper-max-BC, and Clipper-max-aBH (Section 2.5), the only four Clipper variant algorithms applicable to 1vs1 enrichment analysis. The results in Fig. 2 show that, regardless of the contrast scores being difference (diff) or maximum (max), the BC procedure always guarantees the FDR control under a range of target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$. Notably, in terms of power, the two contrast scores consistently have different advantages under the two background scenarios: Clipper-max-BC has higher power under the homogeneous background, while Clipper-diff-BC is more powerful under the heterogeneous background. Considering that the heterogeneous scenario is prevalent in high-throughput biological data, the difference contrast score is preferred. As the power of Clipper-diff-BC drops when q is too small ($q \leq 3\%$) and d is not too large ($d = 1000$), we consider the aBH procedure as an alternative to control the FDR. The results show that Clipper-diff-aBH is indeed more powerful when Clipper-diff-BC lacks power; however, Clipper-diff-aBH cannot guarantee the exact FDR control as Clipper-diff-BC does. Therefore, Clipper uses **Clipper-diff-BC** by default in 1vs1 enrichment analysis, and it recommends users to increase q when too few discoveries are made; if users reject this option, then Clipper would use Clipper-diff-aBH to increase power for the current q.
- **2vs1 enrichment analysis.** Under each of the 6 settings, we compared Clipper-diff-GZ and Clipper-max-GZ (Section 2.5), the only two Clipper variant algorithms applicable to 2vs1 enrichment analysis. For either algorithm, we further compared two numbers of permutation equivalence classes: $h = 1$ or 2, where the latter is $h_{\max} = \binom{3}{1} - 1$—the maximum number of equivalence classes that do not include the identity permutation. Note that h is a required input parameter for the GZ procedure. The results in Fig. 4 show that, regardless of h and the contrast score definition—maximum (max) or difference (diff), the GZ procedure always guarantees the FDR control under all target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$. In terms of power, Clipper-max-GZ($h = 1$) is consistently more powerful than the other three Clipper variants under all settings. Therefore, Clipper uses **Clipper-max-GZ($h = 1$)** by default in enrichment analysis with unequal number of replicates under two conditions.
- **3vs3 enrichment analysis.** Under each of the 12 settings, we compared five Clipper variant algorithms: Clipper-diff-BC, Clipper-diff-aBH, Clipper-max-BC, Clipper-max-aBH, and Clipper-max-GZ (Section 2.5). Fig. 6 shows the comparison of the first four variants: regardless of

the contrast scores being difference (diff) or maximum (max), the BC procedure simultaneously guarantees the FDR control and achieves good power under a range of target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$. Similar to the results in the 1vs1 enrichment analysis, Clipper-max-BC has higher power under the homogeneous background, while Clipper-diff-BC is more powerful under the heterogeneous background. By the same reasoning—the prevalent heterogeneous scenarios in high-throughput biological data—we prefer the difference contrast score. Unlike the 1vs1 enrichment analysis, here Clipper-diff-BC is consistently as powerful as Clipper-diff-aBH, even when q is small, but Clipper-diff-aBH cannot guarantee the exact FDR control. Therefore, Clipper-diff-BC achieves the overall best performance among the first four Clipper variants. Given that the GZ procedure is also applicable to this setting, we further compared Clipper-diff-BC with Clipper-max-GZ($h = 1$), the most powerful Clipper variant with the GZ procedure and the default Clipper implementation in the 2vs1 enrichment and differential analyses and the 3vs3 differential analysis. The results in Supplementary Fig. S5 show that while both **Clipper-diff-BC** and Clipper-max-GZ($h = 1$) control the FDR, the former is more powerful. Hence, we will use Clipper-diff-BC as the default when both conditions have more than one and the same number of replicates.

- **2vs1 differential analysis.** Similar to 2vs1 enrichment analysis, under each of the 6 settings, we compared Clipper-diff-GZ and Clipper-max-GZ (Section 2.5) with $h = 1$ or 2. The results in Fig. 4 show that, regardless of h and the contrast score definition—maximum (max) or difference (diff), the GZ procedure always guarantees the FDR control under a range of target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$. Notably, in terms of power, Clipper-diff-GZ($h = 2$) is the most powerful when when q is very small ($q \leq 2\%$) under Poisson and negative binomial settings, while Clipper-max-GZ($h = 1$) is the most powerful otherwise. Considering that Clipper-max-GZ($h = 1$) outperforms the other three Clipper variants in most cases, Clipper uses **Clipper-max-GZ(**$h = 1$**)** by default in 2vs1 differential analysis, and it recommends users to use Clipper-diff-GZ($h = 2$) when too few discoveries are made.

- **3vs3 differential analysis.** Under each of the 12 settings, we compared Clipper-diff-GZ, and Clipper-max-GZ (Section 2.5) with $h = 1$, 3 or 9, where $h = 9$ is $h_{\max} = \binom{6}{3}/2 - 1$—the maximum number of equivalence classes that do not include the identity permutation. The results in Fig. 8 show that, regardless of h and the contrast score definition—maximum (max) or difference (diff), the GZ procedure always guarantees the FDR control under a range of target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$. In terms of power, Clipper-max-GZ($h = 1$) is consistently more powerful than the other Clipper variant algorithms under all settings. Therefore, Clipper uses **Clipper-max-GZ(**$h = 1$**)** by default in 3vs3 differential analysis.

In summary, whenever Clipper-diff-BC is applicable (enrichment analysis with equal number of replicates under two conditions), it is chosen as the default Clipper implementation; otherwise, Clipper-max-GZ($h = 1$) is the default.

3.3 Comparison of Clipper with generic FDR control methods

For each design and analysis, we compared the default Clipper implementation with other generic FDR control methods. Specifically, seven generic methods (BH-pool, qvalue-pool, BH-pair-mis, qvalue-pair-mis, BH-pair-2as1, qvalue-pair-2as1, and locfdr-emp) are included in all designs and analyses. The two methods relying on correct model specification, BH-pair-correct and qvalue-pair-correct, are only included in the 3vs3 enrichment and differential analyses, because it is almost impossible to correctly specify a model with fewer than three replicates per condition. The permutation-based method, locfdr-swap, is excluded from the 1vs1 enrichment analysis because it requires at least one condition to have

more than one replicate. The comparison results in Figs. 3, 5, 7, 9 and Figs. S1, S2 , S3, S4 show that Clipper controls the FDR and is overall more powerful than other methods, excluding those that fail to control the FDR, under all settings. Clipper is also shown to be more robust to the number of features and the existence of outliers than other methods.

Among the generic methods, the two p-value-based methods using the pooled approach—BH-pool and qvalue-pool—are the most powerful under the homogeneous background, an ideal scenario where features follow the same distribution under the background condition. However, they cannot control the FDR under the heterogeneous background, a realistic scenario in biological studies. Besides, they cannot control the FDR when the number of features is small (Fig. 3). These results confirm that the pooled approach relies on the homogeneous background assumption and a large number of features. The four p-value-based methods using the paired approach—BH-pair-mis, qvalue-pair-mis, BH-pair-2as1, and qvalue-pair-2as1—fail to control the FDR by a large margin in most cases, and rarely when they control the FDR, they lack power. These results confirm that the paired approach relies on correct model specification to control the FDR; however, correct model specification is hardly achievable with no more than three replicates per condition. Even when models are correctly specified (an idealistic scenario), the two paired methods—BH-pair-correct and qvalue-pair-correct, though can control the FDR as expected, fail in the existence of outliers (Fig. 7). The two local-fdr-based methods—locfdr-emp and locfdr-swap—achieve FDR control under all designs and analyses; however, they are less powerful than Clipper in most cases.

In detail, in both enrichment analyses (1vs1, 2vs1, and 3vs3 designs) and differential analyses (2vs1 and 3vs3 designs), Clipper consistently controls the FDR, and it outperforms the generic methods in most cases with few exceptions under the realistic, heterogeneous background, where features do not follow the same distribution under the background condition. Under the idealistic homogeneous background, Clipper is still powerful and only second to BH-pool and qvalue-pool, which, however, cannot control the FDR under the heterogeneous background. The few exceptions where Clipper is not the most powerful under the heterogeneous background include: in 2vs1 differential analysis, Clipper is second to locfdr-emp when the target FDR threshold is small ($q \leq 3\%$) and data are generated from the Poisson or negative binomial distribution (Fig. 5), a case where the default Clipper can be replaced by Clippder-diff-GZ($h = 2$) for better power.

Compared to the two "oracle" methods (BH-pair-correct and qvalue-pair-correct) that use the paired approach and the correct model specification, which cannot be achieved except in simulations, Clipper still demonstrates strong performance: it is only outperformed by the two oracle methods in 3vs3 differential analysis when data are generated from the Poisson distribution (Fig. 9). It is worth noting that correct model specification is unrealistic for the paired approach when the distribution involves unknown parameters in addition to the mean parameter. For example, in 3vs3 differential analysis, we find that valid FDR control for BH-pair-correct and qvalue-pair-correct requires the correct specification of the dispersion parameter value, which is unknown in practice, in the negative binomial distribution; otherwise, correct specification of the distribution family (i.e., negative binomial distribution) alone cannot guarantee the FDR control (Fig. S6).

4 Omics data applications

We demonstrate the use of Clipper in four omics data applications: peak calling from ChIP-seq data, peptide identification from MS data, DEG identification from RNA-seq data, and DIR identification from Hi-C data. The first two applications are enrichment analyses, and the last two are differential analyses. In each application, we compare Clipper with mainstream bioinformatics methods to demonstrate

Clipper's superiority in FDR control and detection power.

4.1 Peak calling from ChIP-seq data (enrichment analysis I)

ChIP-seq is a genome-wide experimental assay for measuring binding sites of a DNA-associated protein [41], often a transcription factor that activates or represses gene expression [42, 43]. ChIP-seq data are crucial for studying gene expression regulation, and the key analysis is to identify genomic regions with enriched sequence reads in ChIP-seq data. These regions are likely to be bound by the targeted protein and thus of biological interest. The identification of these regions is termed "peak calling" in ChIP-seq data analysis.

As the identified peaks are subject to experimental validation that is often expensive [44], it is essential to control the FDR of peak identification to reduce unnecessary costs. The two most highly-cited peak-calling methods are MACS2 [1] and Homer [2], both of which claim to control the FDR for their identified peaks. Specifically, both MACS2 and HOMER assume that the read counts for each putative peak (one count per sample/replicate) follow the Poisson distribution, and they use the paired approach (with biologically-driven modifications) to assign each putative peak a p-value and a corresponding Storey's q-value. Then given a target FDR threshold $0 < q < 1$, they call the putative peaks with q-values $\leq q$ as identified peaks. Despite being popular, MACS2 and HOMER have not been verified for their FDR control capacity in the literature to our knowledge.

To verify the FDR control of MACS2 and HOMER, we used ENCODE ChIP-seq data of cell line GM12878 [45] and generated synthetic data with known true spiked-in peaks to evaluate the FDR and power. We examined the actual FDR and power in a range of target FDR thresholds: $q \in \{1\%, 2\%, \cdots, 10\%\}$. Fig. 10a shows that MACS2 and Homer cannot control the FDR as standalone peak-calling methods. However, with Clipper as an add-on, both MACS2 and Homer can guarantee the FDR control (see Section S4 for the detailed implementation of Clipper in this peak calling analysis). This result demonstrates the flexibility and usability of Clipper for reducing false discoveries in bioinformatics analysis.

Technically, the failed FDR control by MACS2 and HOMER is attributable to the likely model misspecification and test mis-formulation in their use of the paired approach. Both MACS2 and HOMER assume the Poisson distribution for sequence read counts in a putative peak; however, it has been widely acknowledged that read counts are over-dispersed and thus better modeled by the negative binomial distribution [46]. In addition, MACS2 uses one-sample tests to compute p-values when two-sample tests should have been performed. As a result, the p-values of MACS2 and Homer are questionable, so using their p-values for FDR control would not lead to success. (Note that MACS2 does not use p-values to control the FDR but instead swaps experimental and background samples to calculate the empirical FDR; yet, we emphasize that controlling the empirical FDR does not guarantee the FDR control.) As a remedy, Clipper strengthens both methods to control the FDR while maintaining high power.

4.2 Peptide identification from MS data (enrichment analysis II)

The state-of-the-art proteomics studies use MS experiments and database search algorithms to identify and quantify proteins in biological samples. In a typical proteomics experiment, a protein mixture sample is first digested into peptides and then measured by tandem MS as mass spectra, which encode peptide sequence information. The process that decodes mass spectra is called "peptide identification," which converts mass spectra into peptide sequences in a protein sequence database via search algorithms. The search process matches each mass spectrum to peptide sequences in the database and outputs

the best match, called a "peptide-spectrum match" (PSM). The identified PSMs are used to identify and quantify proteins in a high-throughput manner.

False PSMs could occur when mass spectra are matched to wrong peptide sequences due to issues such as low-quality spectra, data-processing errors, and incomplete protein databases, causing problems in the downstream protein identification and quantification [47]. Therefore, a common goal of database search algorithms is to simultaneously control the FDR and maximize the number of identified PSMs, so as to maximize the number of proteins identified in a proteomics study [3, 48, 49]. A widely used FDR control strategy is the target-decoy search, where mass spectra of interest are matched to peptide sequences in both the original (target) database and a decoy database, which contains artificial false protein sequences. The resulting PSMs are called the target PSMs and decoy PSMs, respectively. The decoy PSMs, i.e., matched mass spectrum and decoy peptide pairs, are known to be false and thus used by database search algorithms to control the FDR. Mainstream database search algorithms output a q-value for each PSM, target or decoy. The target PSMs whose q-values are no greater than q, the target FDR threshold, will be reported as discoveries.

We used the first comprehensive benchmark dataset from an archaea species *Pyrococcus furiosus* (Pfu) to examine the FDR control and power of a popular database search algorithm Mascot [3]. Using this benchmark dataset, we demonstrate that, as an add-on, Clipper improves the power of Mascot. Specifically, Clipper treats mass spectra as features. For each mass spectrum, Clipper considers its measurement under the experimental condition as the $-\log_{10}$-transformed q-value of the target PSM that includes it, and its measurement under the background condition as the $-\log_{10}$-transformed q-value of the decoy PSM that includes it. Then Clipper decides which mass spectra and their corresponding target PSMs are discoveries. Based on the benchmark dataset, we examined the empirical FDR, i.e., the FDP calculated based on the true positives and negatives, and the power of Mascot with or without Clipper as an add-on, for a range of target FDR thresholds: $q \in \{1\%, 2\%, \cdots, 10\%\}$. Fig. 10b shows that although Mascot and Mascot+Clipper both control the FDR, Mascot+Clipper consistently improves the power, thus enhancing the peptide identification efficiency of proteomics experiments.

4.3 DEG identification from RNA-seq data (differential analysis I)

RNA-seq data measure genome-wide gene expression levels in biological samples. An important use of RNA-seq data is the DEG analysis, which aims to discover genes whose expression levels change between two conditions. The FDR is a widely used criterion in DEG analysis [4–9].

We compared Clipper with two popular DEG identification methods DESeq2 [5] and edgeR [4]. Specifically, we applied Clipper to two gene expression matrices, one per condition, to identify DEGs. To verify the FDR control of Clipper, DESeq2 and edgeR, we generated realistic synthetic data from real RNA-seq data of classical and non-classical human monocytes [50], i.e., two conditions (see Section S4 for details about the generation of synthetic data). With ground truths (true DEGs and non-DEGs), the synthetic data allow the evaluation of the FDR and power of each method for a range of target FDR thresholds: $q \in \{1\%, 2\%, \cdots, 10\%\}$. Fig. 10c shows that Clipper consistently controls the FDR and achieves high power. In contrast, DESeq2 fails to control the FDR; edgeR controls the FDR but has lower power than Clipper does.

IHW is a popular procedure for boosting the power of p-value-based FDR control methods by incorporating feature covariates [27]. We used IHW as an add-on for DESeq2 and edgeR by adding every gene's mean expression across replicates and conditions as that gene's covariate, as suggested in [51]. The result in Fig. 10c shows that DESeq2+IHW and edgeR+IHW have similar performance in FDR control and power to DESeq2 and edgeR, respectively.

4.4 DIR analysis of Hi-C data (differential analysis II)

Hi-C experiments are widely used to investigate spatial organizations of chromosomes and map chromatin interactions across the genome. A Hi-C dataset is often processed and summarized into an interaction matrix, whose rows and columns represent manually binned chromosomal regions and whose (i, j)-th entry represents the measured contact intensity between the i-th and j-th binned regions. The DIR analysis aims to identify pairs of genomic regions whose interaction intensities differ between conditions. Same as DEG analysis, DIR analysis also uses the FDR as a decision criterion [11–13].

We compared Clipper with three popular DIR identification methods: diffHic [13], FIND [12] and multiHiCcompare [11]. Specifically, we applied Clipper to DIR identification by treating pairs of genomic regions as features and interaction intensities as measurements. To verify the FDR control of Clipper, diffHiC, FIND and multiHiCcompare, we generated realistic synthetic data from real interaction matrices of ENCODE cell line GM12878 [45] with true spiked-in DIRs to evaluate the FDR and power (see Section S4 for details about the generation of synthetic data). We examined the actual FDR and power in a range of target FDR thresholds: $q \in \{1\%, 2\%, \cdots, 10\%\}$. Fig. 10d shows that Clipper and diffHic are the only two methods that consistently control the FDR, while multiHiCcompare and FIND fail by a large margin. In terms of power, Clipper outperforms diffHic except for $q = 0.01$ and 0.02, even though Clipper has not been optimized for Hi-C data analysis. This result demonstrates Clipper's general applicability and strong potential for DIR analysis.

5 Discussion

In this paper, we proposed a new statistical framework, Clipper, for identifying interesting features with FDR control from high-throughput data. Clipper avoids the use of p-values and makes the control of FDR more reliable and flexible. We used comprehensive simulation studies to verify the FDR control by Clipper under various settings. We demonstrate that Clipper outperforms existing generic FDR control methods in terms of power and robustness to model misspecification. We further applied Clipper to four popular bioinformatics analyses: peak calling from ChIP-seq data, peptide identification from MS data, DEG identification from RNA-seq data, and DIR identification from Hi-C data. Our results indicate that Clipper can provide a powerful add-on to existing bioinformatics tools to improve the reliability of FDR control and thus the reproducibility of scientific discoveries.

Clipper is a flexible framework that is easily generalizable to identify a variety of interesting features. The core component of Clipper summarizes each feature's measurements under each condition into an informative statistic (e.g., the sample mean); then Clipper combines each feature's informative statistics under two conditions into a contrast score to enable FDR control. The current implementation of Clipper only uses the sample mean as the informative statistic to identify the interesting features that have distinct expected values under two conditions. However, by modifying the informative statistic, we can generalize Clipper to identify the features that are interesting in other aspects, e.g., having different variances between two conditions. Regarding the contrast score, Clipper makes careful choices between two contrast scores, difference and maximum, based on the number of replicates and the analysis task (enrichment vs. differential). Future studies are needed to explore other contrast scores and their power with respect to data characteristics and analysis tasks.

We have demonstrated the broad application potential of Clipper in various bioinformatics data analyses. Specifically, when used as an add-on to established, popular bioinformatics methods such as MACS2 for peak calling and Mascot for peptide identification, Clipper guaranteed the desired FDR control and in some cases boosted the power. However, many more careful thoughts are needed to escalate Clipper into standalone bioinformatics methods for specific data analyses, for which data pro-

cessing and characteristics (e.g., peak lengths, GC contents, proportions of zeros, and batch effects) must be properly accounted for before Clipper is used for the FDR control [52, 53]. We expect that the Clipper framework will propel future development of bioinformatics methods by providing a flexible p-value-free approach to control the FDR, thus improving the reliability of scientific discoveries.

In this work, we validated the FDR control by Clipper using extensive and concrete simulations including both model-based and real-data-based data generation with ground truths. In contrast, in most bioinformatics method papers, the FDR control was merely mentioned but rarely validated. Many of them assumed that using the BH procedure on p-values would lead to valid FDR control; however, the reality is often otherwise because p-values would be invalid when model assumptions were violated or the p-value calculation was problematic. Here we voice the importance of validating the FDR control in bioinformatics method development, and we use this work as a demonstration. We believe that Clipper provides a powerful booster to this movement. As a p-value-free alternative to the classic p-value-based BH procedure, Clipper relies less on model assumptions and is thus more robust to model misspecifications, making it an appealing choice for FDR control in diverse high-throughput biomedical data analyses.

6 Software

The `Clipper` `R` package is available at https://github.com/JSB-UCLA/Clipper/.

Acknowledgements

The authors would like to thank Dr. Yu-Cheng Yang for his suggestions on the figures and R package. The authors also appreciate the comments and feedback from Tianyi Sun, Kexin Li, and other members of the Junction of Statistics and Biology at UCLA (http://jsb.ucla.edu).

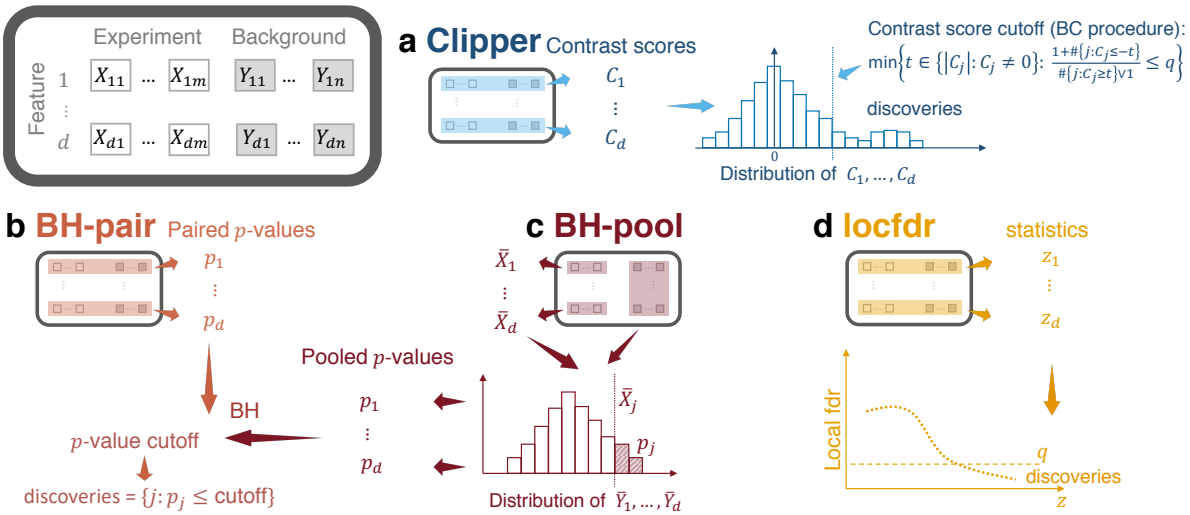

Figure 1: Illustration of Clipper and three generic FDR control approaches: BH-pair, BH-pool, and locfdr. The input data are d features with m and n repeated measurements under the experimental and background conditions, respectively. **(a)** Clipper computes a contrast score for each feature based on the feature's m and n measurements, decides a contrast-score cutoff, and calls the features with contrast scores above the cutoff as discoveries. **(b)** BH-pair computes a p-value for each feature based on the feature's m and n measurements, sets a p-value cutoff, and calls the features with p-values below the cutoff as discoveries. **(c)** BH-pool constructs a null distribution from the d features averaged (across n replicates) measurements under the background condition, calculate a p-value for each feature based on the null distribution and the feature's averaged (across m replicates) measurements under the experimental condition, sets a p-value cutoff, and calls the features with p-values below the cutoff as discoveries. **(d)** locfdr computes a summary statistic for each feature based on the feature's m and n measurements, estimates the empirical null distribution and the empirical distribution of the statistic across features, computes a local fdr for each feature, sets a local fdr cutoff, and calls the features with local fdr below the cutoff as discoveries.

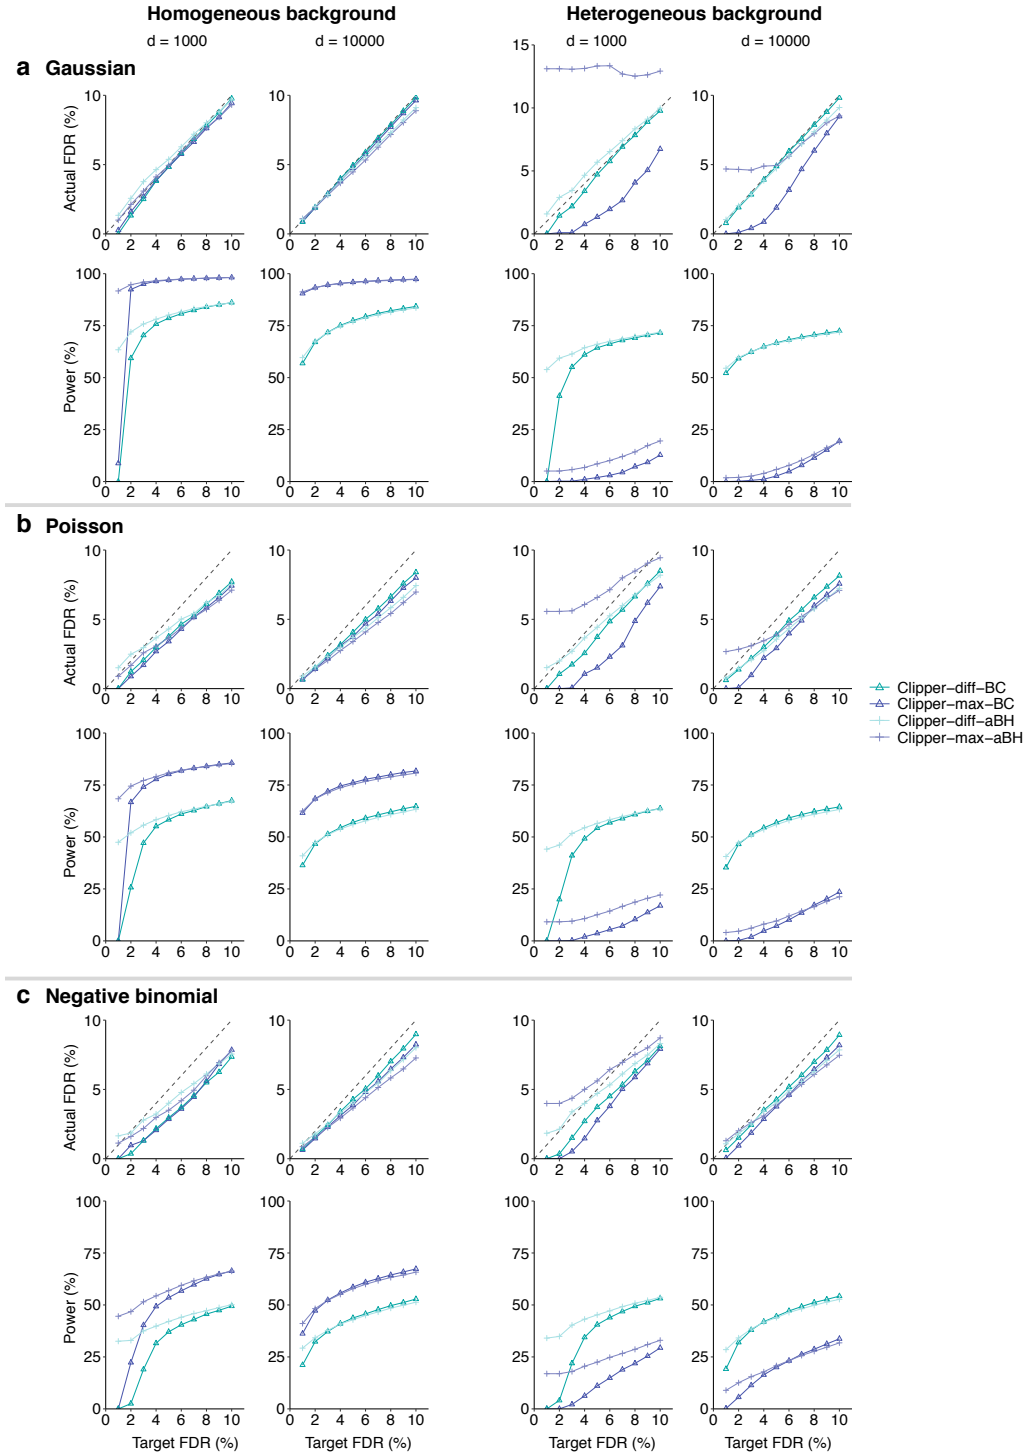

Figure 2: In 1vs1 enrichment analysis, comparison of four Clipper variant algorithms (Clipper-diff-BC, Clipper-diff-aBH, Clipper-max-BC, and Clipper-max-aBH) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with $d = 1000$ or 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Clipper-diff-BC is chosen as the default implementation under this scenario.

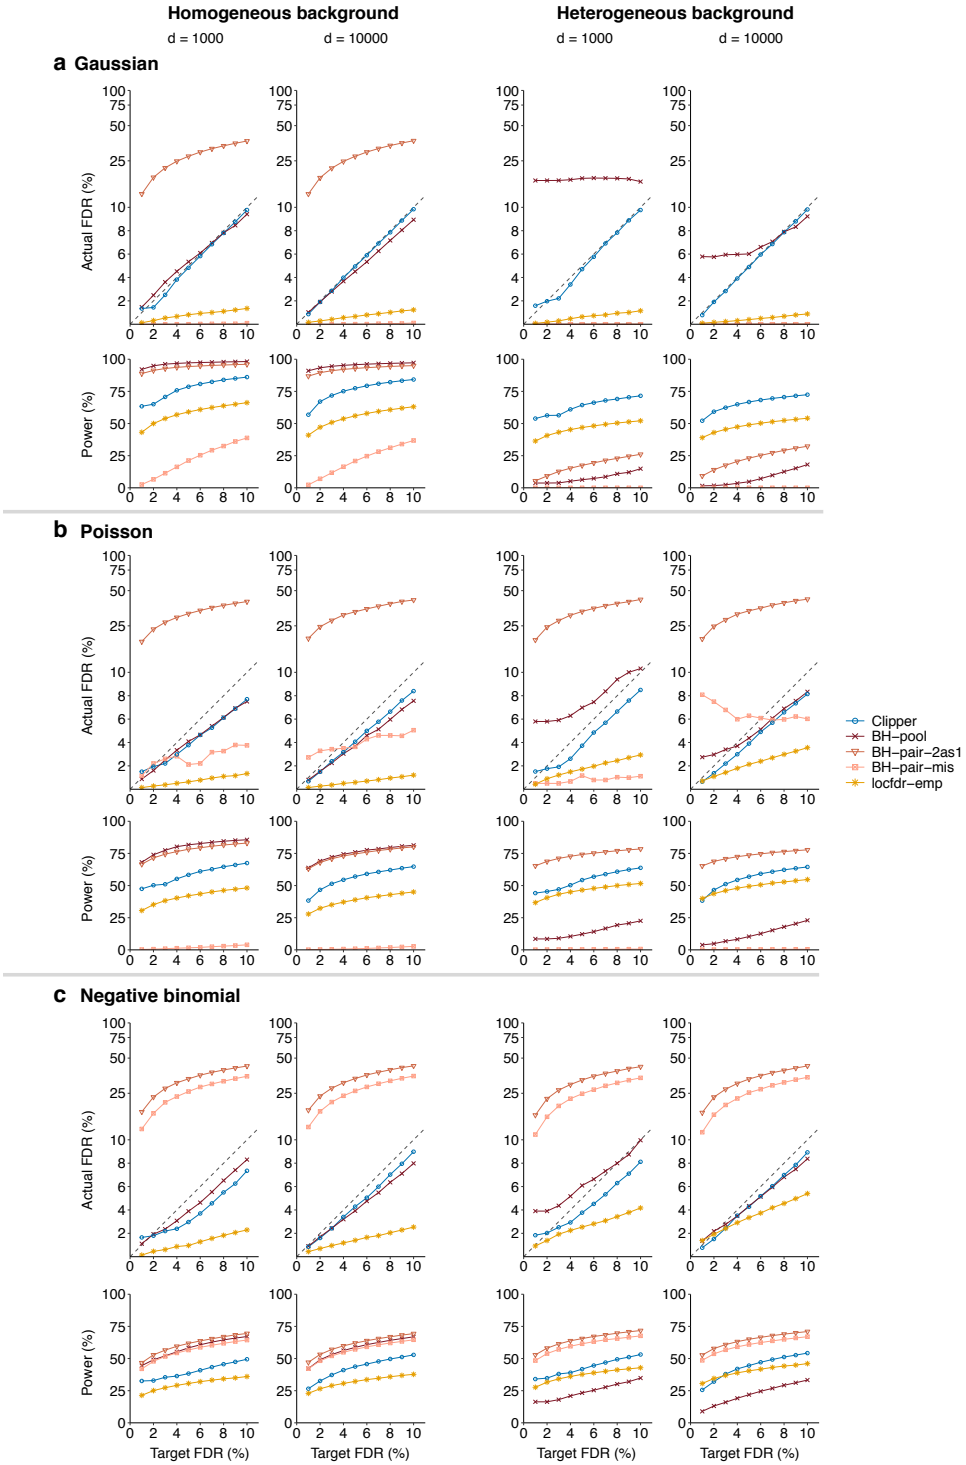

Figure 3: In 1vs1 enrichment analysis, comparison of Clipper and four other generic FDR control methods (BH-pool, BH-pair-2as1, BH-pair-mis, and locfdr-emp) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with $d = 1000$ or 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background.

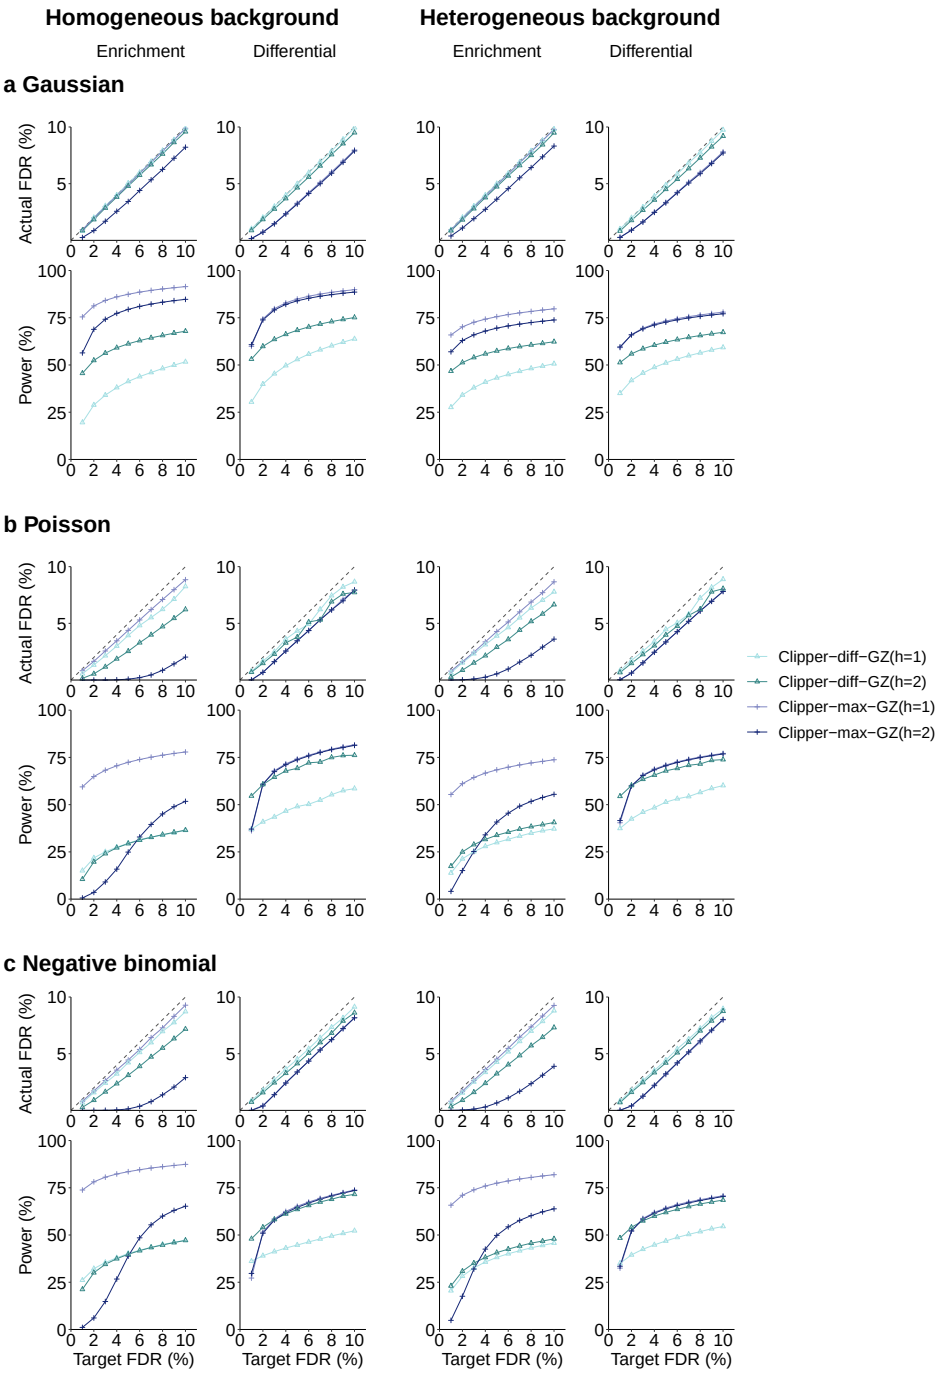

Figure 4: In 2vs1 enrichment analysis (columns 1 and 3) and differential analysis (columns 2 and 4), comparison of four Clipper variant algorithms (Clipper-diff-GZ(h=1), Clipper-diff-GZ(h=2), Clipper-max-GZ(h=1), Clipper-max-GZ(h=2)) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous(two left columns) and heterogeneous (two right columns) background scenarios. Clipper-max-GZ(h=1) is chosen as the default implementation under this scenario.

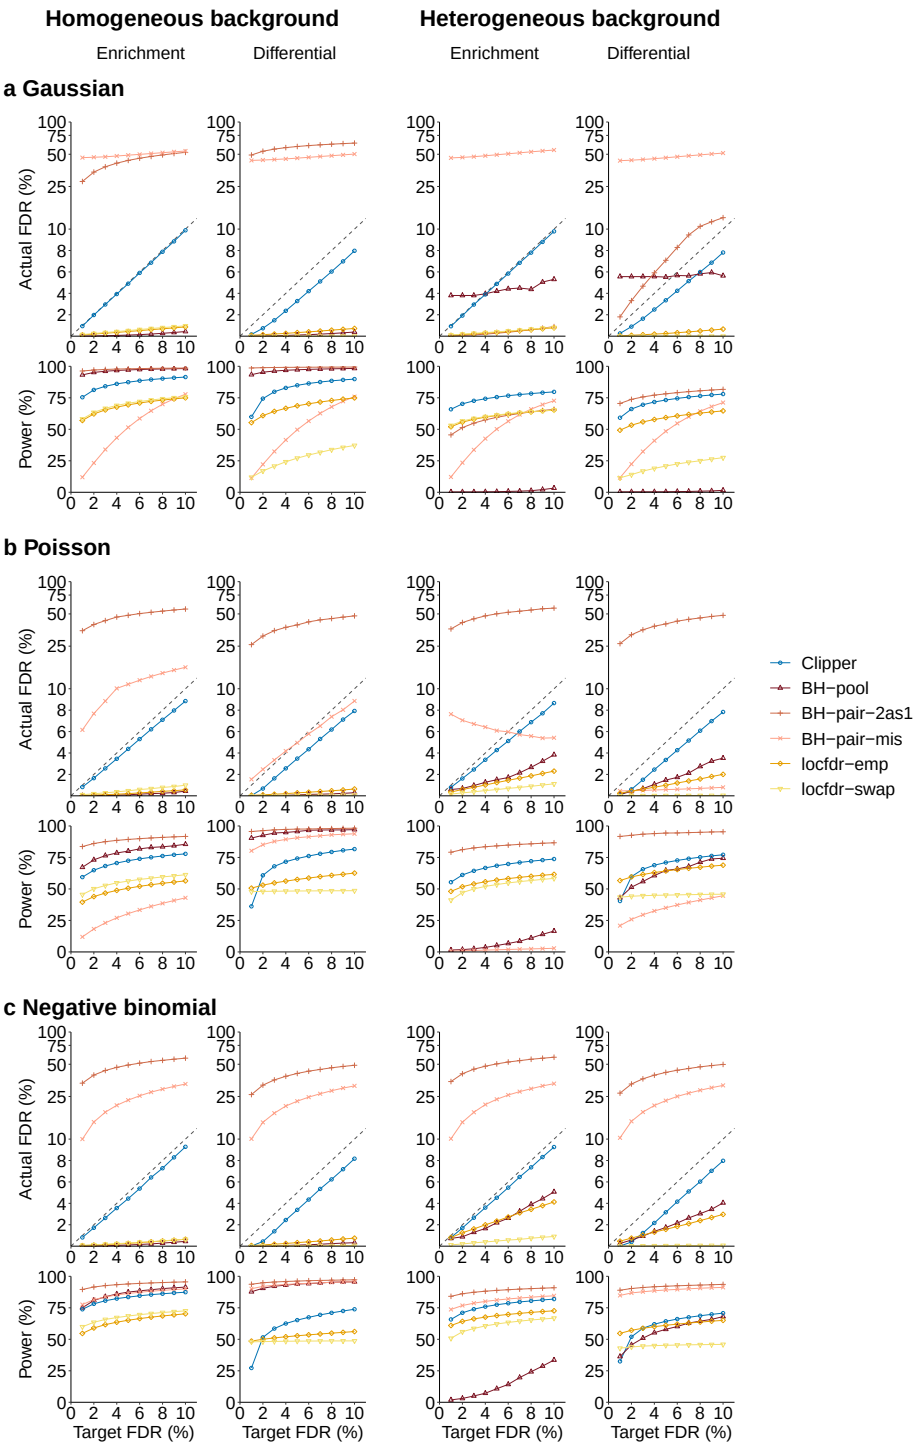

Figure 5: In 2vs1 enrichment analysis (columns 1 and 3) and differential analysis (columns 2 and 4), comparison of Clipper and five other generic FDR control methods (BH-pooled, BH-pair-2as1, BH-pair-mis, locfdr-emp, and locfdr-swap) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous(two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background (except for differential analysis with $q \leq 2\%$).

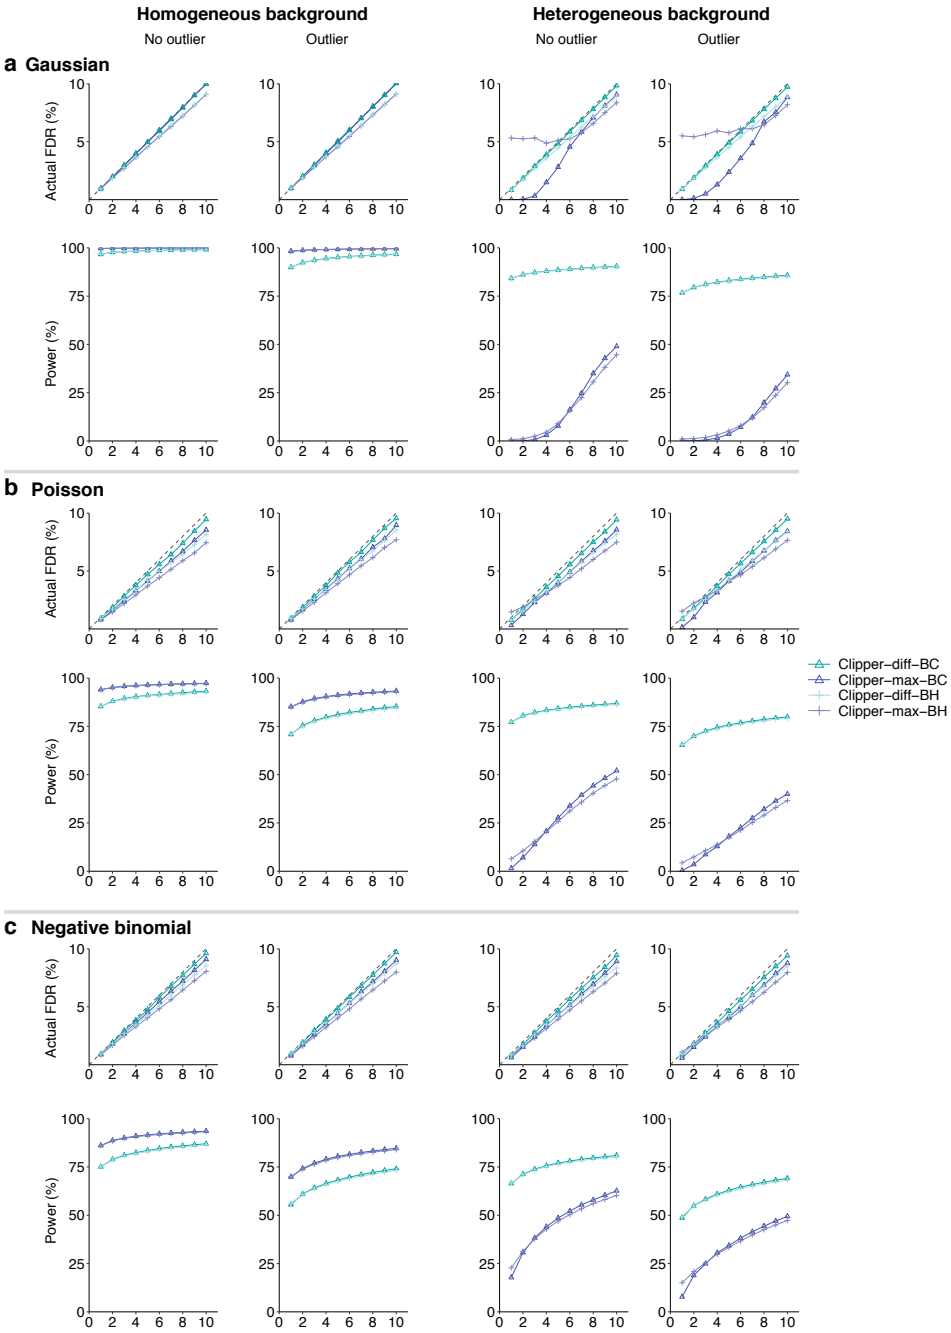

Figure 6: In 3vs3 enrichment analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of four Clipper variant algorithms (Clipper-diff-BC, Clipper-diff-aBH, Clipper-max-BC, and Clipper-max-aBH) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Clipper-diff-BC is chosen as the default implementation under this scenario.

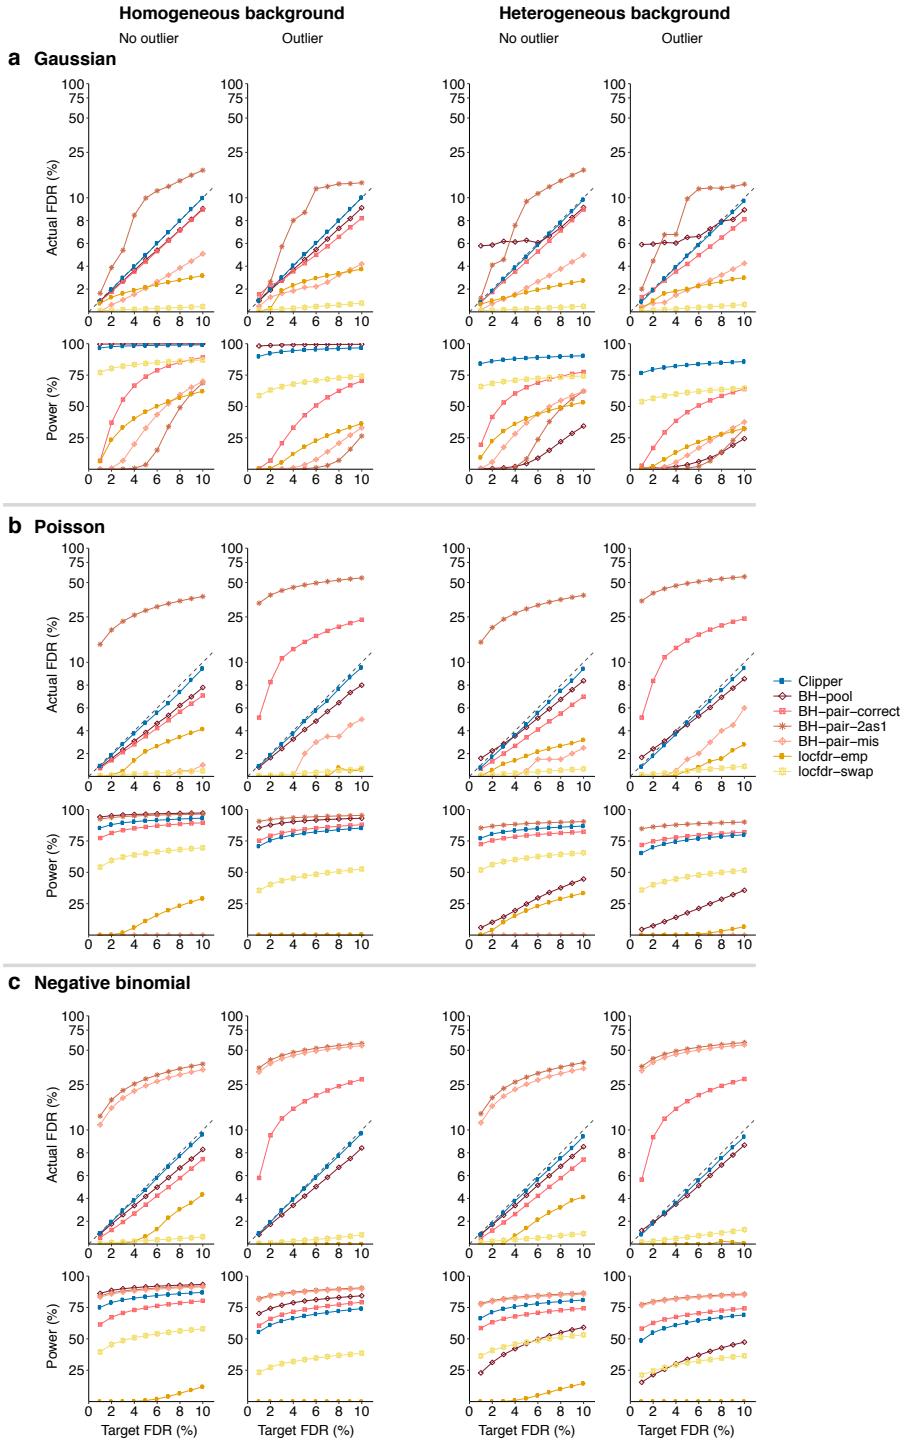

Figure 7: In 3vs3 enrichment analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of Clipper and six other generic FDR control methods (BH-pooled, BH-pair-correct, BH-pair-2as1, BH-pair-mis, locfdr-emp, and locfdr-swap)in terms of their FDR control and power in 3vs3 enrichment analysis with possible outliers. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background.

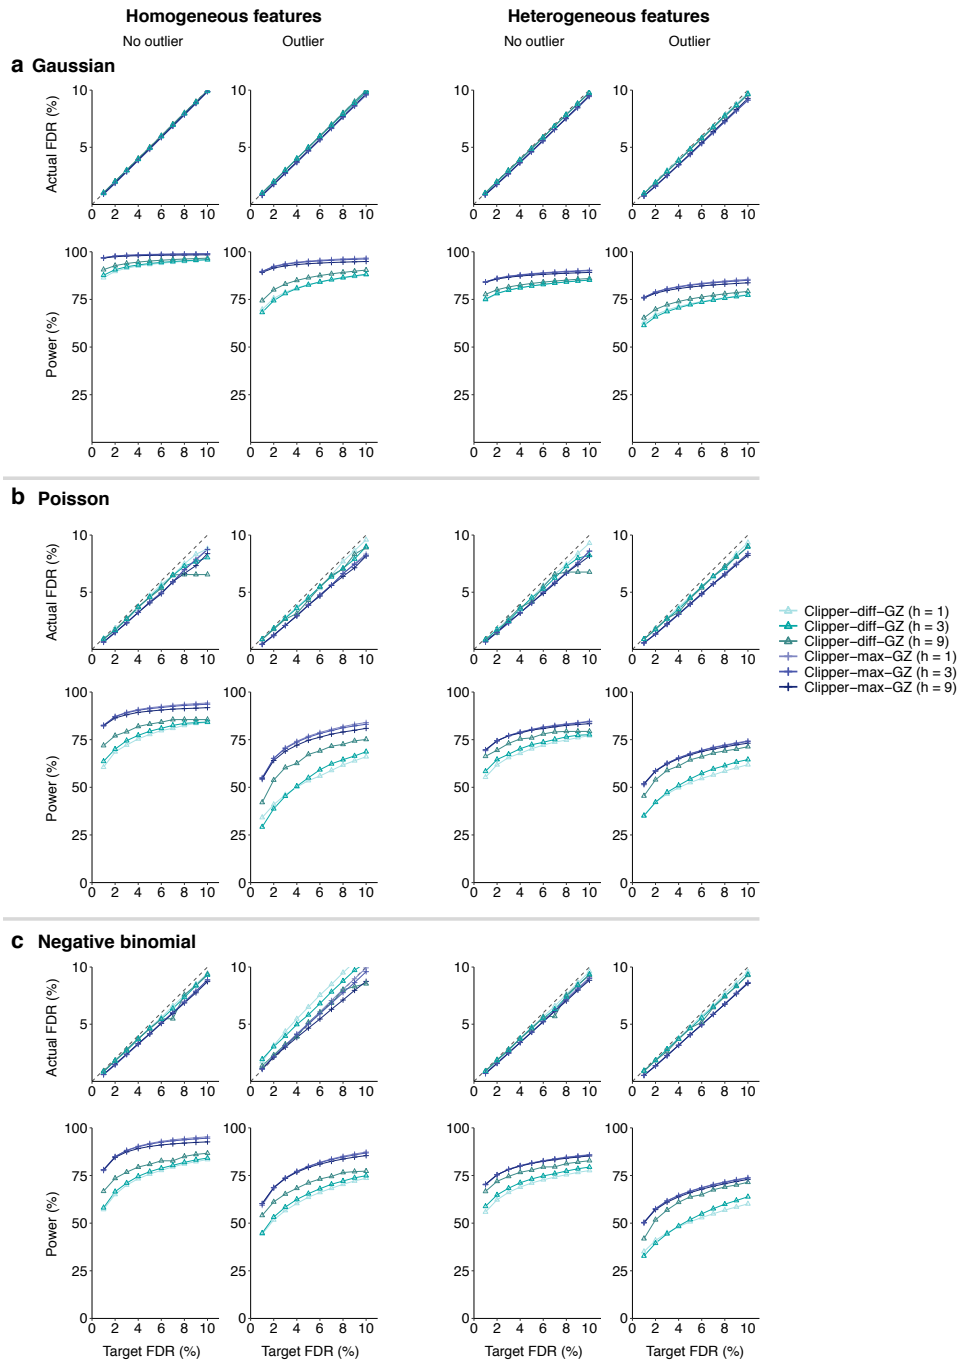

Figure 8: In 3vs3 differential analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of six Clipper variant algorithms (Clipper-diff-GZ(h=1), Clipper-diff-GZ(h=3), Clipper-diff-GZ(h=9), Clipper-max-GZ(h=1), Clipper-max-GZ(h=3), and Clipper-max-GZ(h=9)) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Clipper-max-GZ(h=1) is chosen as the default implementation under this scenario.

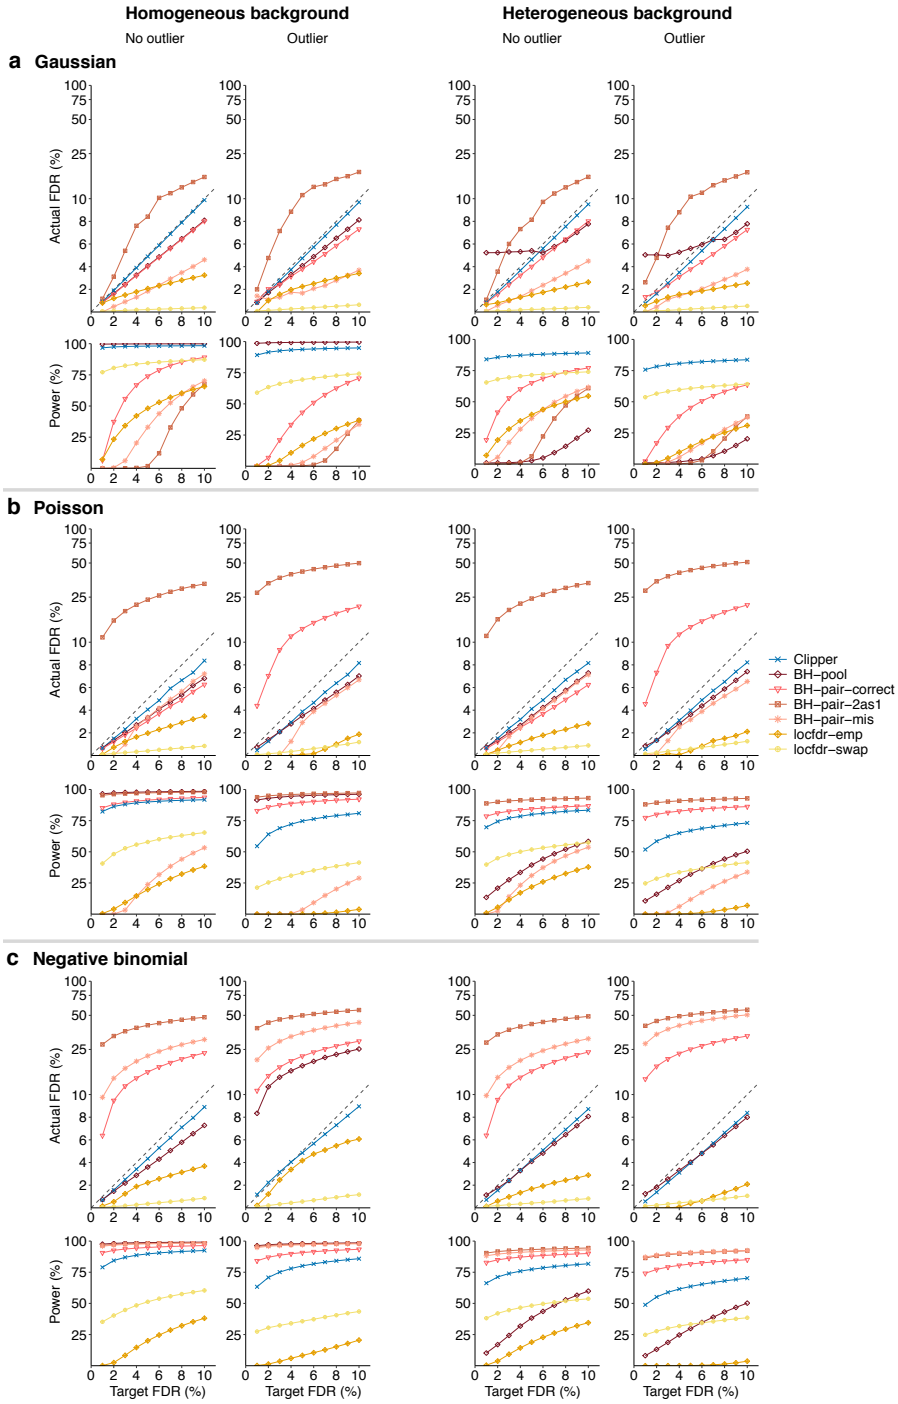

Figure 9: In 3vs3 differential analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of Clipper and six other generic FDR control methods (BH-pooled, BH-pair-correct, BH-pair-2as1, BH-pair-mis, locfdr-emp, and locfdr-swap) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background (except for Poisson distribution where Clipper is second to BH-pair-correct, an idealistic method).

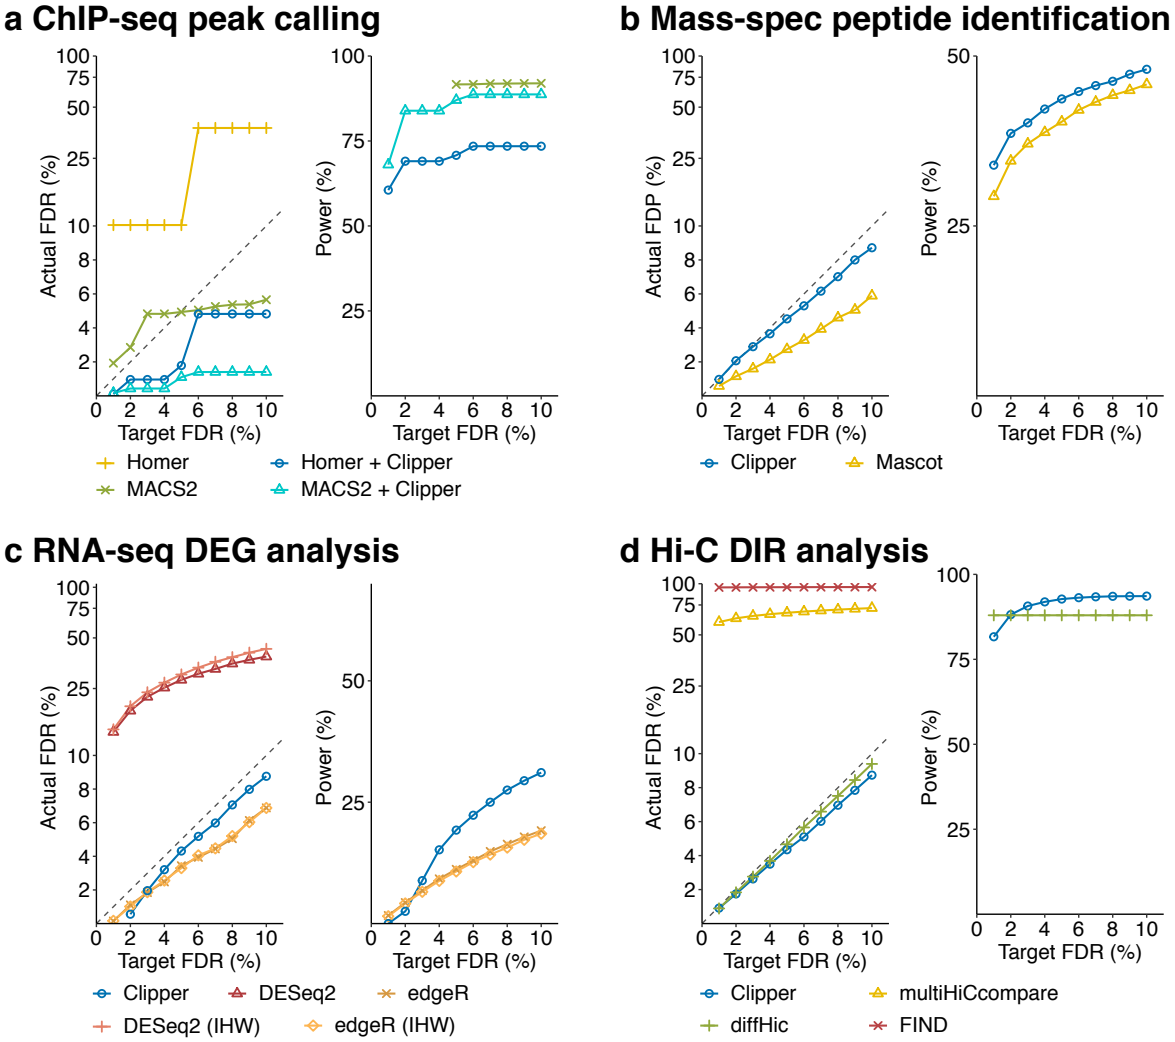

Figure 10: Comparison of Clipper and popular bioinformatics methods in terms of FDR control and power in **(a)** peaking calling analysis on synthetic ChIP-seq data, **(b)** peptide identification on real proteomics data, **(c)** DEG analysis on synthetic bulk RNA-seq data, and **(d)** DIR analysis on synthetic Hi-C data . In all four panels, target FDR levels range from 0.01 to 0.1. Points above the dashed line indicate failed FDR control; when this happens, the power of the corresponding methods is not shown, including Homer in (a), MACS2 for target FDR less than 5% in (a), DESeq2 and DESeq2 (IHW) in (c), and multiHICcompare and FIND in (d). In all four applications, Clipper controls the FDR while maintaining high power, demonstrating Clipper's broad applicability in high-throughput data analyses.

Supplementary Material

S1 Existing generic FDR control methods

To facilitate our discussion, we re-introduce the notations for data. For feature $j = 1, \ldots, d$, we use $\boldsymbol{X}_j = (X_{j1}, \ldots, X_{jm})^\top \in \mathbb{R}^m_{\geq 0}$ and $\boldsymbol{Y}_j = (Y_{j1}, \ldots, Y_{jn})^\top \in \mathbb{R}^n_{\geq 0}$ to denote its measurements under the experimental and background conditions, respectively. We assume that $X_{j1}, \ldots, X_{jm}$ are identically distributed, so are $Y_{j1}, \ldots, Y_{jn}$. Let $\mu_{Xj} = \mathbb{E}[X_{j1}]$ and $\mu_{Yj} = \mathbb{E}[Y_{j1}]$ denote the expected measurement of feature j under the two conditions, respectively. Then we use $\bar{X}_j$ to denote the sample average of $X_{j1}, \cdots, X_{jm}$ and $\bar{Y}_j$ the sample average of $Y_{j1}, \cdots, Y_{jn}$.

S1.1 P-value-based methods

Here we describe the details of p-value-based FDR control methods, including BH-pair, BH-pool, qvalue-pair, and qvalue-pool. Each of these four methods first computes p-values using either the pooled approach or the paired approach, and it then relies on the BH procedure [14] or Storey's qvalue procedure [25] for FDR control. In short, every p-value-based method is a combination of a p-value calculation approach and a p-value thresholding procedure. Below we introduce two p-value calculation approaches (paired and pooled) and two p-value thresholding procedures (BH and Storey's qvalue).

S1.1.1 P-value calculation approaches

The paired approach. The paired approach examines one feature at a time and compares its measurements between two conditions. Besides the ideal implementation, i.e., the *correct paired approach* that uses the correct model to calculate p-values, we also include commonly-used flawed implementations that either misspecify the distribution, i.e., the *misspecified paired approach*, or mis-formulate the two-sample test as a one-sample test, i.e., the *2as1 paired approach*.

Here we use the negative binomial distribution as an example to demonstrate the ideas of the correct, misspecified, and 2as1 paried approaches. Suppose that for each feature j, its measurements under each condition follow a negative binomial distribution, and the two distributions under the two conditions have the same dispersion; that is, $X_{j1}, \cdots, X_{jm} \overset{\text{i.i.d.}}{\sim} \mathrm{NB}\left(\mu_{Xj}, \theta_j\right)$; $Y_{j1}, \cdots, Y_{jn} \overset{\text{i.i.d.}}{\sim} \mathrm{NB}\left(\mu_{Yj}, \theta_j\right)$, where θ_j is the dispersion parameter such that the variance $\mathrm{Var}(X_{ji}) = \mu_{Xj} + \theta_j \mu_{Xj}^2$.

- The correct paired approach assumes that the two negative binomial distributions have the same dispersion parameter θ_j, and it uses the two-sample test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$ (enrichment analysis) or $H_1 : \mu_{Xj} \neq \mu_{Yj}$ (differential analysis).

- The misspecified paired approach misspecifies the negative binomial distribution as Poisson, and it uses the two-sample test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$ (enrichment analysis) or $H_1 : \mu_{Xj} \neq \mu_{Yj}$ (differential analysis).

- The 2as1 paired approach bluntly assumes $\mu_{Yj} = \bar{Y}_j$, and it performs the one-sample test based on $X_{j1}, \ldots, X_{jm}$ for the null hypotheses $H_0 : \mu_{Xj} = \bar{Y}_j$ against the alternative hypothesis $H_1 : \mu_{Xj} > \bar{Y}_j$ (enrichment analysis) or $H_1 : \mu_{Xj} \neq \bar{Y}_j$ (differential analysis).

The pooled approach. The pooled approach pools all features' average measurements under the background condition $\left\{\bar{Y}_j\right\}_{j=1}^d$ to form a null distribution, and it calculates a p-value for each feature j

by comparing $\bar{X}_j$ to the null distribution. Specifically, in enrichment analysis, the p-value of feature j is computed as:

$$p_j = \frac{\operatorname{card}\left(\left\{k : \bar{Y}_k \geq \bar{X}_j\right\}\right)}{d}\,.$$

In differential analysis, the p-value of feature j is computed as:

$$p_j = 2 \cdot \min\left(\frac{\operatorname{card}\left(\left\{k : \bar{Y}_k \geq \bar{X}_j\right\}\right)}{d}\,,\ \frac{\operatorname{card}\left(\left\{k : \bar{Y}_k \leq \bar{X}_j\right\}\right)}{d}\right)\,.$$

S1.1.2 P-value thresholding procedures for FDR control

Definition 5 (BH procedure for thresholding p-values [14]) *The features' p-values $p_1, \ldots, p_d$ are sorted in an ascending order $p_{(1)} \leq p_{(2)} \leq \ldots \leq p_{(d)}$. Given the target FDR threshold q, the Benjamini–Hochberg (BH) procedure finds a p-value cutoff T^{BH} as*

$$T^{\mathrm{BH}} := p_{(k)},\ \textit{where}\ k = \max\left\{j = 1, \ldots, d : p_{(j)} \leq \frac{j}{d} q\right\}. \tag{13}$$

Then BH outputs $\left\{j : p_j \leq T^{\mathrm{BH}}\right\}$ as discoveries.

Definition 6 (Storey's qvalue procedure for thresholding p-values [25]) *The features' p-values $p_1, \ldots, p_d$ are sorted in an ascending order $p_{(1)} \leq p_{(2)} \leq \ldots \leq p_{(d)}$. Let $\hat{\pi}_0$ denote an estimate of the probability $P(\textit{the } i\textit{-th feature is uninteresting})$ (see Storey [25] for details). Storey's qvalue procedure defines the q-value for $p_{(d)}$ as*

$$\hat{q}(p_{(d)}) := \frac{\hat{\pi}_0 \cdot d \cdot p_{(d)}}{\operatorname{card}\left(\left\{k : p_k \leq p_{(d)}\right\}\right)} = \hat{\pi}_0 \cdot p_{(d)}\,.$$

Then for $j = d-1, d-2, \ldots, 1$, the q-value for $p_{(j)}$ is defined as:

$$\hat{q}(p_{(j)}) := \min\left(\hat{q}(p_{(j+1)}),\ \frac{\hat{\pi}_0 \cdot d \cdot p_{(j)}}{\operatorname{card}\left(\left\{k : p_k \leq p_{(j)}\right\}\right)}\right).$$

Then Storey's qvalue procedure outputs $\{j : \hat{q}(p_j) \leq q\}$ as discoveries.

We use function `qvalue` from `R` package `qvalue` (with default estimate $\hat{\pi}_0$) to calculate q-values.

S1.2 Local-fdr-based methods

The FDR is statistical criterion that ensures the reliability of discoveries as a whole. In contrast, the local fdr focuses on the reliability of each discovery. The definition of the local fdr relies on some pre-computed summary statistics z_j for feature j, $j = 1, \ldots, d$. In the calculation of local fdr, $\{z_1, \ldots, z_d\}$ are assumed to be realizations of an abstract random variable Z that represents any feature. Let p_0 or p_1 denote the prior probability that any feature is uninteresting or interesting, with $p_0 + p_1 = 1$. Let $f_0(z) := \mathbb{P}(Z = z \,|\, \text{uninteresting})$ or $f_1(z) := \mathbb{P}(Z = z \,|\, \text{interesting})$ denote the conditional probability density of Z at z given that Z represents an uninteresting or interesting feature. Thus by Bayes' theorem, the posterior probability of any feature being uninteresting given its summary statistic $Z = z$ is

$$\mathbb{P}(\text{uninteresting} \mid Z = z) = p_0 f_0(z)/f(z)\,, \tag{14}$$

where $f(z) := p_0 f_0(z) + p_1 f_1(z)$ is the marginal probability density of Z. Accordingly, the local fdr of feature j is defined as follows.

Definition 7 (Local fdr [54]) *Given notations defined above, the local fdr of feature j is defined as*

$$\textit{local-fdr}_j := f_0(z_j)/f(z_j)\,.$$

Because $p_0 \leq 1$, local-fdr$_j$ is an upper bound of the posterior probability of feature j being uninteresting given its summary statistic z_j, defined in (14)*.*

Note that another definition of the local fdr is the posterior probability $\mathbb{P}(\text{uninteresting} \mid z)$ in (14) [16]. Although this other definition is more reasonable, we do not use it but choose Definition 7 because the estimation of p_0 is ususally difficult. Another reason is that uninteresting features are the dominant majority in high-throughput biological data, so p_0 is often close to 1. Hence, Definition 7, though being more stringent, guarantees the FDR control and would not lead to much power loss.

We define local-fdr-based methods as a type of FDR control methods by thresholding local fdrs of features under the target FDR threshold q. Although the local fdr is different from FDR, it has been shown that thresholding the local fdrs at q will approximately control the FDR under q [54]. This makes local-fdr-based methods competitors against Clipper and p-value-based methods.

Every local-fdr-based method is a combination of a local fdr calculation approach and a local fdr thresholding procedure. Below we introduce two local fdr calculation approaches (empirical null and swapping) and one local fdr thresholding procedure. After the combination, we have two local-fdr-based methods: locfdr-emp and locfdr-swap.

S1.2.1 Local fdr calculation approaches

With $z_1, \ldots, z_d$, the calculation of local fdr defined in Definition 7 requires the estimation of f_0 and f, two probability densities. f is estimated by nonparametric density estimation, and f_0 is estimated by either the empirical null approach [54] or the swapping approach, which shuffles replicates between conditions [16]. With the estimated $\hat{f}$ and $\hat{f}_0$, the estimated local fdr of feature j is

$$\widehat{\textit{local-fdr}}_j := \hat{f}_0(z_j)/\hat{f}(z_j)\,. \tag{15}$$

The empirical null approach. This approach assumes a parametric distribution, typically the Gaussian distribution, to estimate f_0. Then with the density estimate $\hat{f}$, the local fdr is estimated for each feature j. The implementation of this approach depends on the numbers of replicates.

- In 1vs1 enrichment and differential analyses, we define z_j as

$$z_j := \frac{D_j}{\sqrt{\frac{1}{d}\sum_{j=1}^{d}\left(D_j - \bar{D}\right)^2}}\,,$$

  where $D_j = X_{j1} - Y_{j1}$ and $\bar{D} = \sum_{j=1}^{d} D_j/d$.

- In 2vs1 enrichment and differential analyses, we define z_j as

$$z_j := \frac{\bar{X}_j - Y_{j1}}{\sqrt{\frac{s_{X_j}^2}{2}}}\,,$$

  where $s_{X_j}^2 = \sum_{i=1}^{2}(X_{ji} - \bar{X}_j)^2$.

- In $m$vs$n$ enrichment and differential analyses with $m, n \geq 2$, we define z_j as the two-sample

t-statistic with unequal variances:

$$z_j := \frac{\bar{X}_j - \bar{Y}_j}{\sqrt{\frac{s^2_{Xj}}{m} + \frac{s^2_{Yj}}{n}}},$$

where $s^2_{Xj} = \frac{1}{m-1}\sum_{i=1}^{m}(X_{ji} - \bar{X}_j)^2$ and $s^2_{Yj} = \frac{1}{n-1}\sum_{i=1}^{n}(Y_{ji} - \bar{Y}_j)^2$ are the sample variances of feature j under the experimental and background conditions.

Then $\{\widehat{\text{locfdr}}_j\}_{j=1}^d$ are estimated from $\{z_j\}_{j=1}^d$ by function `locfdr` in `R` package `locfdr` (with default arguments).

The swapping approach. This approach swaps $\lceil m/2 \rceil$ replicates under the experimental condition with $\lceil n/2 \rceil$ replicates under the background condition. Then it calculates the summary statistic for each feature on the swapped data, obtaining $z'_1, \ldots, z'_d$. Finally, it estimates f_0 and f by applying kernel density estimation to $z'_1, \ldots, z'_d$ and $z_1, \ldots, z_d$, respectively (by function `kde` in `R` package `ks`). With $\hat{f}_0$ and $\hat{f}$, $\{\widehat{\text{locfdr}}_j\}_{j=1}^d$ are calculated by Definition 7.

The implementation of this approach depends on the numbers of replicates. Below are three special cases included in this work.

- In 1vs1 enrichment and differential analyses, the swapping approach is inapplicable because interesting features stay as interesting after the swapping.
- In 2vs1 enrichment and differential analyses, we define z_j and z'_j as

$$z_j = \frac{X_{j1} + X_{j2}}{2} - Y_{j1}\,,$$
$$z'_j = \frac{X_{j1} + Y_{j1}}{2} - X_{j2}\,.$$

- In 3vs3 enrichment and differential analyses with, we define z_j and z'_j as

$$z_j = \frac{X_{j1} + X_{j2}}{2} - \frac{Y_{j1} + Y_{j2}}{2}\,,$$
$$z'_j = \frac{X_{j1} + Y_{j1}}{2} - \frac{X_{j2} + Y_{j2}}{2}\,.$$

Then we apply kernel density estimation to $\{z_j\}_{j=1}^d$ and $\{z'_j\}_{j=1}^d$ to obtain $\hat{f}$ and $\hat{f}_0$, respectively. By (15), we calculate $\{\widehat{\text{locfdr}}_j\}_{j=1}^d$.

S1.2.2 The local fdr thresholding procedure

Definition 8 (locfdr procedure) *Given the local fdr estimates $\{\widehat{\text{local-fdr}}_j\}_{j=1}^d$ and the target FDR threshold q, the locfdr procedure outputs $\{j = 1, \ldots, d : \widehat{\text{local-fdr}}_j \le q\}$ as discoveries.*

S2 Data generation and detailed implementation of the paired approach (a p-value calculation approach) in simulation studies

We describe how we simulated data and how we implemented the paired approach in different simulation settings: 1vs1 enrichment analysis, 2vs1 enrichment analysis, 3vs3 enrichment analysis, 2vs1 differential analysis, and 3vs3 differential analysis, combined with three distribution families (Gaussian, Poisson, and negative binomial) and two background scenarios (homogeneous and heterogeneous). Under some settings, we considered different numbers of features and the existence of outliers.

In each simulation setting, we generated 200 simulated datasets, computed an FDP and an empirical power on each dataset, and averaged the 200 FDPs and 200 empirical powers to approximate the FDR and power, repsectively. For notation simplicity, we use $N(\mu, \sigma^2)$ to denote the Gaussian distribution with mean μ and variance σ^2, $\mathrm{Pois}(\lambda)$ to denote the Poisson distribution with mean λ, and $\mathrm{NB}(\mu, \theta)$ to denote the negative binomial distribution with mean μ and dispersion θ (such that its variance equals $\mu + \theta\mu^2$).

S2.1 1vs1 enrichment analysis

We simulated data with $d = 1000$ and $10,000$ features under two background scenarios and three distributional families—a total of 12 settings. In each setting, 10% of the features are interesting ($\mu_{Xj} > \mu_{Yj}$), and the rest are uninteresting (with $\mu_{Xj} = \mu_{Yj}$). Recall that $\mathcal{N}$ denotes the set of uninteresting features.

Gaussian distribution

We simulated data from Gaussian using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 0$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 0$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(5, 1)$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $N(0, 2)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(5, 1)$.
- We independently generated X_{j1} from $N(\mu_{Xj}, 1)$ and Y_{j1} from $N(\mu_{Yj}, 1)$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that the null distribution of $X_{j1} - Y_{j1}$, $j = 1, \ldots, d$ is $N(0, \hat{\sigma}^2)$, where

$$\hat{\sigma}^2 = \frac{1}{d-1} \sum_{j=1}^{d} \left(X_{j1} - \frac{1}{d} \sum_{j=1}^{d} X_{j1} \right)^2 + \frac{1}{d-1} \sum_{j=1}^{d} \left(Y_{j1} - \frac{1}{d} \sum_{j=1}^{d} Y_{j1} \right)^2 .$$

This is a misspecified model that assumes that μ_{Xj}'s are all equal and so are μ_{Yj}'s. Then we computed the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $X_{j1} - Y_{j1}$ in $N(0, \hat{\sigma}^2)$, i.e., $1 - \Phi\left(\frac{X_{j1} - Y_{j1}}{\hat{\sigma}}\right)$, where Φ is the cumulative distribution function of $N(0, 1)$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $N(Y_{j1}, 1)$ conditioning on the observed Y_{j1} as the null distribution of X_{j1}. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the right tail probability of X_{j1} in $N(Y_{j1}, 1)$, i.e., $1 - \Phi\left(X_{j1} - Y_{j1}\right)$.

Poisson distribution

We simulated data from Poisson using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{Pois}(40)$.

- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $\mathrm{Pois}(20)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{Pois}(40)$.

- We independently generated X_{j1} from $\mathrm{Pois}(\mu_{Xj})$ and Y_{j1} from $\mathrm{Pois}(\mu_{Yj})$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we first defined a log-transformation $f(x) = \log(x + 0.01)$, which we applied to X_{j1} and Y_{j1}, $j = 1, \ldots, d$. We assumed that the null distribution of $f(X_{j1}) - f(Y_{j1})$, $j = 1, \ldots, d$ is $N(0, \hat{\sigma}^2)$, where

$$\hat{\sigma}^2 = \frac{1}{d-1} \sum_{j=1}^{d} \left(f(X_{j1}) - \frac{1}{d} \sum_{j=1}^{d} f(X_{j1}) \right)^2 + \frac{1}{d-1} \sum_{j=1}^{d} \left(f(Y_{j1}) - \frac{1}{d} \sum_{j=1}^{d} f(Y_{j1}) \right)^2 .$$

This model misspecifies the Poisson distribution as the log-normal distribution.

Then we computed the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $f(X_{j1}) - f(Y_{j1})$ in $N(0, \hat{\sigma}^2)$, i.e., $1 - \Phi\left(\frac{f(X_{j1}) - f(Y_{j1})}{\hat{\sigma}}\right)$, where Φ is the cumulative distribution function of $N(0, 1)$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $\mathrm{Pois}(Y_{j1})$ conditioning on the observed Y_{j1} as the null distribution of X_{j1}. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the right tail probability of X_{j1} in $\mathrm{Pois}(Y_{j1})$, i.e., $\mathbb{P}(Z \geq X_{j1})$ where $Z \sim \mathrm{Pois}(Y_{j1})$.

Negative binomial distribution

We simulated data from negative binomial using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{NB}(45, 45^{-1})$.

- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $\mathrm{NB}(20, 20^{-1})$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{NB}(45, 45^{-1})$.

- We independently generated X_{j1} from $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ and Y_{j1} from $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that for each uninteresting feature j, Y_{j1} and X_{j1} follow the same Poisson distribution. We calculated the p-value of feature j from a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$ using function `poisson.test` in R package `stats`.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $\mathrm{NB}(Y_{j1}, Y_{j1}^{-1})$ conditioning on the observed Y_{j1} as the null distribution of X_{j1}. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the right tail probability of X_{j1} in $\mathrm{NB}(Y_{j1}, Y_{j1}^{-1})$.

S2.2 2vs1 enrichment analysis

We simulated data with $d = 10,000$ features under two background scenarios and three distributional families—a total of 6 settings. In each setting, 10% of the features are interesting ($\mu_{Xj} > \mu_{Yj}$) and the rest are uninteresting (with $\mu_{Xj} = \mu_{Yj}$). Recall that $\mathcal{N}$ denotes the set of uninteresting features.

Gaussian distribution

We simulated data from Gaussian using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 0$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 0$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(5, 1)$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j \in \mathcal{N}}$ i.i.d. from $N(0, 2)$ and set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. We next generated $\{\mu_{Yj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(0, 2)$ and $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(5, 1)$.
- We independently generated X_{j1} and X_{j2} from $N(\mu_{Xj}, 1)$ and Y_{j1} from $N(\mu_{Yj}, 1)$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that the null distribution of $\frac{1}{2}(X_{j1} + X_{j2}) - Y_{j1}$, $j = 1, \ldots, d$ is $N(0, \hat{\sigma}^2)$, where

$$\hat{\sigma}^2 = \frac{1}{2(2d-1)} \sum_{j=1}^{d} \sum_{i=1}^{2} \left(X_{ji} - \frac{1}{2d} \sum_{j=1}^{d} \sum_{i=1}^{2} X_{ji} \right)^2 + \frac{1}{d-1} \sum_{j=1}^{d} \left(Y_{j1} - \frac{1}{d} \sum_{j=1}^{d} Y_{j1} \right)^2 .$$

This is a misspecified model that assumes μ_{Xj}'s are all equal and so are μ_{Yj}'s. Then we computed the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $\frac{1}{2}(X_{j1} + X_{j2}) - Y_{j1}$ in $N(0, \hat{\sigma}^2)$, i.e., $1 - \Phi\left(\frac{\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}}{\hat{\sigma}}\right)$, where Φ is the cumulative distribution function of $N(0, 1)$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $N(Y_{j1}, 1/2)$ conditioning on the observed Y_{j1} as the null distribution of $\frac{1}{2}(X_{j1} + X_{j2})$. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $\frac{1}{2}(X_{j1} + X_{j2})$ in $N(Y_{j1}, 1/2)$, i.e., $1 - \Phi\left(\frac{\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}}{1/\sqrt{2}}\right)$.

Poisson distribution

We simulated data from Poisson using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{Pois}(40)$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^{d}$ i.i.d. from $\mathrm{Pois}(20)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{Pois}(40)$.
- We independently generated X_{j1} and X_{j2} from $\mathrm{Pois}(\mu_{Xj})$ and Y_{j1} from $\mathrm{Pois}(\mu_{Yj})$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we first defined a log-transformation $f(x) = \log(x + 0.01)$, which we applied to X_{j1} and Y_{j1}, $j = 1, \ldots, d$. We assumed that the null distribution of $f(X_{j1}) + f(X_{j2}) - 2f(Y_{j1})$, $j = 1, \ldots, d$ is $N(0, \hat{\sigma}^2)$, where

$$\hat{\sigma}^2 = \frac{6}{d-1} \sum_{j=1}^{d} \left(f(Y_{j1}) - \frac{1}{d} \sum_{j=1}^{d} f(Y_{j1}) \right)^2 .$$

This model misspecifies the Poisson distribution as the log-normal distribution.

Then we computed the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $f(X_{j1}) + f(X_{j2}) - 2f(Y_{j1})$ in $N(0, \hat{\sigma}^2)$, i.e., $1 - \Phi\left(\frac{f(X_{j1})+f(X_{j2})-2f(Y_{j1})}{\hat{\sigma}}\right)$, where Φ is the cumulative distribution function of $N(0, 1)$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we assumed that for each uninteresting feature j, X_{j1} and X_{j2} independently follow $\mathrm{Pois}(Y_{j1})$ conditioning on the observed Y_{j1}. Then we calculated the p-value of feature $j = 1, \ldots, d$ by performing a one-sample Poisson test using the R function `poisson.test` for the null hypothesis $H_0 : \mu_{X_j} = Y_{j1}$ against the alternative hypothesis $H_1 : \mu_{X_j} > Y_{j1}$.

Negative binomial distribution

We simulated data from negative binomial using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{NB}(45, 45^{-1})$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^{d}$ i.i.d. from $\mathrm{NB}(20, 20^{-1})$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{NB}(45, 45^{-1})$.
- We independently generated X_{j1} and X_{j2} from $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ and Y_{j1} from $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that for each uninteresting feature j, X_{ji}, $i = 1, 2$ and Y_{j1} follow the same Poisson distribution. We calculated the p-value of feature j from a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$ using the function `poisson.test` in R package `stats`.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $\mathrm{NB}(2Y_{j1}, (2Y_{j1})^{-1})$ conditioning on the observed Y_{j1} as the null distribution of $X_{j1} + X_{j2}$. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $X_{j1} + X_{j2}$ in $\mathrm{NB}(2Y_{j1}, (2Y_{j1})^{-1})$.

S2.3 3vs3 enrichment analysis

We simulated data with and without outliers under two background scenarios and three distributional families—a total of 12 settings. In each setting, we generated $d = 10,000$ features, among which 10% are interesting (with $\mu_{Xj} > \mu_{Yj}$) and the rest are uninteresting (with $\mu_{Xj} = \mu_{Yj}$).

Under the settings with outliers, we generated $\{O_{ji}^X : j = 1, \ldots, d; i = 1, \ldots, 3\}$ and $\{O_{ji}^Y : j = 1, \ldots, d; i = 1, \ldots, 3\}$ i.i.d. from $\mathrm{Bernoulli}(0.1)$, where $O_{ji}^X = 1$ or $O_{ji}^Y = 1$ indicates X_{ji} or Y_{ji} is an outlier, respectively. Under settings without outliers, $O_{ji}^X = O_{ij}^Y = 0$ for all $j = 1, \ldots, d; i = 1, \ldots, 3$.

Gaussian distribution

- Under the homogeneous background scenario, we set $\mu_{Yj} = 0$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 0$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(5, 1)$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^{d}$ i.i.d. from $N(0, 2)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $N(5, 1)$.

- We independently generated X_{ji} from $N(\mu_{Xj}, 1)$ if $O^X_{ji} = 0$ or from the top 1% percentile of $N(\mu_{Xj}, 1)$ if $O^X_{ji} = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$. Similarly, we independently generated Y_{ji} from $N(\mu_{Yj}, 1)$ if $O^Y_{ji} = 0$ or from the top 1% percentile of $N(\mu_{Yj}, 1)$ if $O^Y_{ji} = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$.

To implement the correct paired approach (as in BH-pair-correct and qvalue-pair-correct), we calculated the p-value of feature j from a two-sample t-test with equal variance for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we calculated the p-value of feature j from a two-sample t-test with unequal variance for the null hypothesis $H_0 : \mu_{X_j} = \mu_{Y_j}$ against the alternative hypothesis $H_1 : \mu_{X_j} > \mu_{Y_j}$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we assumed that for each uninteresting feature j, X_{ji}, $i = 1, \ldots, 3$ are i.i.d. Gaussian with mean $\bar{Y}_j$ conditioning on the observed $\bar{Y}_j$ and unknown variance. We calculated the p-value of feature j using a one-sample t-test for the null hypothesis $H_0 : \mu_{X_j} = \bar{Y}_j$ against the alternative hypothesis $H_1 : \mu_{X_j} > \bar{Y}_j$.

Poisson distribution

We simulated data from Poisson using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{Pois}(40)$.

- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $\mathrm{Pois}(20)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{Pois}(40)$.

- We independently generated X_{ji} from $\mathrm{Pois}(\mu_{Xj})$ if $O^X_{ji} = 0$ or from the top 1% percentile of $\mathrm{Pois}(\mu_{Xj})$ if $O^X_{ji} = 1$, $j = 1, \ldots, d$, $i = 1, \ldots, 3$. Similarly, we independently generated Y_{ji} from $\mathrm{Pois}(\mu_{Yj})$ if $O^Y_{ji} = 0$ or from the top 1% percentile of $\mathrm{Pois}(\mu_{Yj})$ if $O^Y_{ji} = 1$, $j = 1, \ldots, d$; $1, \ldots, 3$.

To implement the correct paired approach (as in BH-pair-correct and qvalue-pair-correct), we calculated the p-value of feature j by performing a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$ using the function `poisson.test` in R package `stats`.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we first defined a log-transformation $f(x) = \log(x + 0.01)$, which we applied to X_{ji} and Y_{ji}, $j = 1, \ldots, d$; $i = 1, \ldots, 3$. We assumed that for each uninteresting feature j, $\{f(X_{ji})\}_{i=1}^3$ and $\{f(Y_{ji})\}_{i=1}^3$ follow Gaussian distributions with mean $\mu_{f(Xj)}$ and $\mu_{f(Yj)}$, respectively. Then we computed the p-value of feature j using a two-sample equal variance t-test for the null hypothesis $H_0 : \mu_{f(Xj)} = \mu_{f(Yj)}$ against the alternative hypothesis $H_1 : \mu_{f(Xj)} > \mu_{f(Yj)}$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we assumed that for each uninteresting feature j, $\{X_{ji}\}_{i=1}^3$ follow $\mathrm{Pois}(\bar{Y}_j)$ conditioning on the observed $\bar{Y}_j$. We calculated the p-value of feature j by performing a one-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \bar{Y}_j$ against the alternative hypothesis $H_1 : \mu_{Xj} > \bar{Y}_j$ using R function `poisson.test` from package `stats`.

Negative binomial distribution

We simulated data from negative binomial using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{NB}(45, 45^{-1})$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^{d}$ i.i.d. from $\mathrm{NB}(20, 20^{-1})$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For interesting features, we generated $\{\mu_{Xj}\}_{j \notin \mathcal{N}}$ i.i.d. from $\mathrm{NB}(45, 45^{-1})$.
- We independently generated X_{ji} from $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ if $O_{ji}^{X} = 0$ or from the top 1% percentile of $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ if $O_{ji}^{X} = 1$, $j = 1, \ldots, d$, $i = 1, \ldots, 3$. Similarly, we independently generated Y_{ji} from $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$ if $O_{ji}^{Y} = 0$ or from the top 1% percentile of $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$ if $O_{ji}^{Y} = 1$, $j = 1, \ldots, d$, $i = 1, \ldots, 3$.

To implement the correct paired approach (as in BH-pair-correct and qvalue-pair-correct), we performed a two-sample negative binomial test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative $H_1 : \mu_{Xj} > \mu_{Yj}$ using $T_j := \sum_{i=1}^{3} X_{ji} - \sum_{i=1}^{3} Y_{ji}$ as the test statistic. We computed the p-value of feature j as the right tail probability

$$\mathbb{P}(T_j \geq t_j) = \sum_{k_1=0}^{\infty} \sum_{k_2=k_1+t_j}^{\infty} \mathbb{P}\left(\sum_{i=1}^{3} X_{ji} \geq k_2\right) \mathbb{P}\left(\sum_{i=1}^{3} Y_{ji} = k_1\right),$$

where t_j is the realization of T_j, $\mathbb{P}(\sum_{i=1}^{3} X_{ji} \geq k_2)$ and $\mathbb{P}(\sum_{i=1}^{3} Y_{ji} = k_1)$ can be estimated from the null distribution of X_{ji} and Y_{ji}, $j = 1, \ldots, d$; $i = 1, \ldots, 3$. As $\sum_{i=1}^{3} X_{ji}$ and $\sum_{i=1}^{3} Y_{ji}$ follow the same distribution under null, we estimated μ_{Xj} and μ_{Yj} as $\hat{\mu}_{Xj} = \hat{\mu}_{Yj} := (\sum_{i=1}^{3} X_{ji} + \sum_{i=1}^{3} Y_{ji})/6$. Then, we calculated $\mathbb{P}(\sum_{i=1}^{3} X_{ji} \geq k_2)$ and $\mathbb{P}(\sum_{i=1}^{3} Y_{ji} = k_1)$ using the estimated distribution of X_{ji} and Y_{ji} as $\mathrm{NB}(\hat{\mu}_{Xj}, (\hat{\mu}_{Xj})^{-1})$ and $\mathrm{NB}(\hat{\mu}_{Yj}, (\hat{\mu}_{Yj})^{-1})$, respectively, $j = 1, \ldots, d$; $i = 1, \ldots, 3$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that for each uninteresting feature j, $\{X_{ji}\}_{j=1}^{3}$ and $\{Y_{ji}\}_{j=1}^{3}$ follow the same Poisson distribution. We calculated the p-value of feature j from a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} > \mu_{Yj}$ using function `poisson.test` in R package `stats`.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $\mathrm{NB}(\sum_{i=1}^{3} Y_{ji}, (\sum_{i=1}^{3} Y_{ji})^{-1})$ conditioning on the observed $\sum_{i=1}^{3} Y_{ji}$ as the null distribution of $\sum_{i=1}^{3} X_{ji}$. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the right tail probability of $\sum_{i=1}^{3} X_{ji}$ in $\mathrm{NB}(\sum_{i=1}^{3} Y_{ji}, (\sum_{i=1}^{3} Y_{ji})^{-1})$.

S2.4 2vs1 differential analysis

We simulated data with $d = 10,000$ features under two background scenarios and three distributional families—a total of 6 settings. In each setting, we generated $d = 10,000$ features, among which 10% features were "up-regulated features" with $\mu_{Xj} > \mu_{Yj}$ and another 10% were "down-regulated features" with $\mu_{Xj} < \mu_{Yj}$.

Gaussian distribution

We simulated data from Gaussian using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 0$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 0$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $N(5, 1)$. For down-regulated features, generated μ_{Xj} i.i.d. from $N(-5, 1)$.

- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $N(0,2)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $N(5,1)$. For down-regulated features, generated μ_{Xj} i.i.d. from $N(-5,1)$.
- We independently generated X_{j1} and X_{j2} from $N(\mu_{Xj},1)$ and Y_{j1} from $N(\mu_{Yj},1)$, $j=1,\ldots,d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that the null distribution of $\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}$, $j=1,\ldots,d$, is $N(0,\hat{\sigma}^2)$, where

$$\hat{\sigma}^2 = \frac{1}{2(2d-1)}\sum_{j=1}^{d}\sum_{i=1}^{2}\left(X_{ji} - \frac{1}{2d}\sum_{j=1}^{d}\sum_{i=1}^{2}X_{ji}\right)^2 + \frac{1}{d-1}\sum_{j=1}^{d}\left(Y_{j1} - \frac{1}{d}\sum_{j=1}^{d}Y_{j1}\right)^2.$$

This is a misspecified model assuming that μ_{Xj}'s are all equal and so are μ_{Yj}'s. Then we computed the p-value of feature $j=1,\ldots,d$ as the two-sided tail probability of $\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}$ in $N(0,\hat{\sigma}^2)$, i.e., $2\cdot\min\left(1-\Phi\left(\frac{\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}}{\hat{\sigma}}\right),\Phi\left(\frac{\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}}{\hat{\sigma}}\right)\right)$, where Φ is the cumulative distribution function of $N(0,1)$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $N(Y_{j1},1)$ conditioning on the observed Y_{j1} as the null distribution of X_{j1}. Then we calculated the p-value of feature $j=1,\ldots,d$ as the two-sided tail probability of $\frac{1}{2}(X_{j1}+X_{j2})$ in $N(Y_{j1},1/2)$, i.e., $2\cdot\min\left(1-\Phi\left(\frac{\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}}{1/\sqrt{2}}\right),\Phi\left(\frac{\frac{1}{2}(X_{j1}+X_{j2})-Y_{j1}}{1/\sqrt{2}}\right)\right)$.

Poisson distribution

We simulated data from Poisson using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj}=20$ for all d features. For uninteresting features, we set $\mu_{Xj}=\mu_{Yj}=20$ for $j\in\mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(60)$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(5)$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $\mathrm{Pois}(20)$. For uninteresting features, we set $\mu_{Xj}=\mu_{Yj}$ for $j\in\mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(60)$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(5)$.
- We independently generated X_{j1} and X_{j2} from $\mathrm{Pois}(\mu_{Xj})$ and Y_{j1} from $\mathrm{Pois}(\mu_{Yj})$, $j=1,\ldots,d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we first defined a log-transformation $f(x)=\log(x+0.01)$, which we applied to X_{j1} and Y_{j1}, $j=1,\ldots,d$. We assumed that the null distribution of $f(X_{j1})+f(X_{j2})-2f(Y_{j1})$, $j=1,\ldots,d$ is $N(0,\hat{\sigma}^2)$, where

$$\hat{\sigma}^2 = \frac{6}{d-1}\sum_{j=1}^{d}\left(f(Y_{j1}) - \frac{1}{d}\sum_{j=1}^{d}f(Y_{j1})\right)^2.$$

This model misspecifies the Poisson distribution as the log-normal distribution. Then we computed the p-value of feature $j=1,\ldots,d$ as the two-sided tail probability of $f(X_{j1})+f(X_{j2})-2f(Y_{j1})$ in $N(0,\hat{\sigma}^2)$, i.e., $2\cdot\min\left(1-\Phi\left(\frac{f(X_{j1})+f(X_{j2})-2f(Y_{j1})}{\hat{\sigma}}\right),\Phi\left(\frac{f(X_{j1})+f(X_{j2})-2f(Y_{j1})}{\hat{\sigma}}\right)\right)$, where Φ is the cumulative distribution function of $N(0,1)$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we assumed that for each uninteresting feature j, X_{j1} and X_{j2} independently follow $\mathrm{Pois}(Y_{j1})$ conditioning on the observed Y_{j1}. Then we calculated the p-value of feature $j=1,\ldots,d$ by performing a one-sample Poisson test using the R function `poisson.test` for the null hypothesis $H_0: \mu_{X_j}=Y_{j1}$ against the alternative hypothesis $H_1: \mu_{X_j}\neq Y_{j1}$.

Negative binomial distribution

We simulated data from negative binomial using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 30$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 30$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(70, 70^{-1})$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(7, 7^{-1})$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^{d}$ i.i.d. from $\mathrm{NB}(30, 30^{-1})$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(70, 70^{-1})$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(7, 7^{-1})$.
- We independently generated X_{j1} and X_{j2} from $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ and Y_{j1} from $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$, $j = 1, \ldots, d$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that for each uninteresting feature j, X_{j1}, X_{j2}, and Y_{j1} follow the same Poisson distribution. We calculated the p-value of feature j from a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} \neq \mu_{Yj}$ using the function `poisson.test` in R package `stats`.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $\mathrm{NB}(2Y_{j1}, (2Y_{j1})^{-1})$ conditioning on the observed Y_{j1} as the null distribution of $X_{j1}+X_{j2}$. Then we calculated the p-value of feature $j = 1, \ldots, d$ as the two-sided tail probability of $X_{j1}+X_{j2}$ in $\mathrm{NB}(2Y_{j1}, (2Y_{j1})^{-1})$, i.e., twice the smaller of the left-tail and right-tail probabilities.

S2.5 3vs3 differential analysis

We simulated data with or without outliers under two background scenarios and three distributional families—a total of 12 settings. In each setting, we generated $d = 10,000$ features, among which 10% features were "up-regulated features" with $\mu_{Xj} > \mu_{Yj}$ and another 10% were "down-regulated features" with $\mu_{Xj} < \mu_{Yj}$.

Under the settings with outliers, we generated $\{O_{ji}^X : j = 1, \ldots, d; i = 1, \ldots, 3\}$ and $\{O_{ji}^Y : j = 1, \ldots, d; i = 1, \ldots, 3\}$ i.i.d. from $\mathrm{Bernoulli}(0.1)$, where $O_{ji}^X = 1$ or $O_{ji}^Y = 1$ indicates X_{ji} or Y_{ji} is an outlier, respectively. Under settings without outliers, $O_{ji}^X = O_{ij}^Y = 0$ for all $j = 1, \ldots, d; i = 1, \ldots, 3$.

Gaussian distribution

- Under the homogeneous background scenario, we set $\mu_{Yj} = 0$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 0$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $N(5, 1)$. For down-regulated features, generated μ_{Xj} i.i.d. from $N(-5, 1)$.
- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^{d}$ i.i.d. from $N(0, 2)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $N(5, 1)$. For down-regulated features, generated μ_{Xj} i.i.d. from $N(-5, 1)$.
- We independently generated X_{ji} from $N(\mu_{Xj}, 1)$ if $O_{ji}^X = 0$ or from the top 1% percentile of $N(\mu_{Xj}, 1)$ if $O_{ji}^X = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$. Similarly, we independently generated Y_{ji} from $N(\mu_{Yj}, 1)$ if $O_{ji}^Y = 0$ or from the top 1% percentile of $N(\mu_{Yj}, 1)$ if $O_{ji}^Y = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$.

To implement the correct paired approach (as in BH-pair-correct and qvalue-pair-correct), we calculated the p-value of feature j from a two-sample t-test with equal variance for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} \neq \mu_{Yj}$.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we calculated the p-value of feature j from a two-sample t-test with unequal variance for the null hypothesis $H_0 : \mu_{X_j} = \mu_{Y_j}$ against the alternative hypothesis $H_1 : \mu_{X_j} \neq \mu_{Y_j}$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we treated $N(\bar{Y}_j, 1)$ conditioning on observed $\bar{Y}_j$ as the null distribution of X_{ji}, $i = 1, \ldots, 3$. We calculated the p-value of feature j using a one-sample t-test for the null hypothesis $H_0 : \mu_{X_j} = \bar{Y}_j$ against the alternative hypothesis $H_1 : \mu_{X_j} \neq \bar{Y}_j$.

Poisson distribution

We simulated data from Poisson using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 20$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 20$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(40)$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(5)$.

- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $\mathrm{Pois}(20)$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(40)$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{Pois}(5)$.

- We independently generated X_{ji} from $\mathrm{Pois}(\mu_{Xj})$ if $O_{ji}^X = 0$ or from the top 1% percentile of $\mathrm{Pois}(\mu_{Xj})$ if $O_{ji}^X = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$. Similarly, we independently generated Y_{ji} from $\mathrm{Pois}(\mu_{Yj})$ if $O_{ji}^Y = 0$ or from the top 1% percentile of $\mathrm{Pois}(\mu_{Yj})$ if $O_{ji}^Y = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$.

To implement the correct paired approach (as in BH-pair-correct and qvalue-pair-correct), we calculated the p-value of feature j by performing a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} \neq \mu_{Yj}$ using function `poisson.test` in R package `stats`.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we first defined a log-transformation $f(x) = \log(x + 0.01)$, which we applied to X_{ji} and Y_{ji}, $j = 1, \ldots, d$; $i = 1, \ldots, 3$. We assumed that for each uninteresting feature j, $\{f(X_{ji})\}_{i=1}^3$ and $\{f(Y_{ji})\}_{i=1}^3$ follow Gaussian distributions with mean $\mu_{f(Xj)}$ and $\mu_{f(Yj)}$, respectively. Then we computed the p-value of feature j using a two-sample equal variance t-test for the null hypothesis $H_0 : \mu_{f(Xj)} = \mu_{f(Yj)}$ against the alternative hypothesis $H_1 : \mu_{f(Xj)} \neq \mu_{f(Yj)}$.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we assumed that for each uninteresting feature j, $\{X_{ji}\}_{i=1}^3$ follow $\mathrm{Pois}(\bar{Y}_j)$ conditioning on the observed $\bar{Y}_j$. We calculated the p-value of feature j by performing a one-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \bar{Y}_j$ against the alternative hypothesis $H_1 : \mu_{Xj} \neq \bar{Y}_j$ using the function `poisson.test` in R package `stats`.

Negative binomial distribution

We simulated data from negative binomial using the following procedure:

- Under the homogeneous background scenario, we set $\mu_{Yj} = 30$ for all d features. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj} = 30$ for $j \in \mathcal{N}$. For up-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(70, 70^{-1})$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(7, 7^{-1})$.

- Under the heterogeneous background scenario, we generated $\{\mu_{Yj}\}_{j=1}^d$ i.i.d. from $\mathrm{NB}(30, 30^{-1})$. For uninteresting features, we set $\mu_{Xj} = \mu_{Yj}$ for $j \in \mathcal{N}$. For up-regulated features, we gener-

ated μ_{Xj} i.i.d. from $\mathrm{NB}(70, 70^{-1})$. For down-regulated features, we generated μ_{Xj} i.i.d. from $\mathrm{NB}(7, 7^{-1})$.

- We independently generated X_{ji} from $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ if $O_{ji}^X = 0$ or from the top 1% percentile of $\mathrm{NB}(\mu_{Xj}, \mu_{Xj}^{-1})$ if $O_{ji}^X = 1$; $j = 1, \ldots, d$, $i = 1, \ldots, 3$. Similarly, we independently generated Y_{ji} from $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$ if $O_{ji}^Y = 0$ or from the top 1% percentile of $\mathrm{NB}(\mu_{Yj}, \mu_{Yj}^{-1})$ if $O_{ji}^Y = 1$, $j = 1, \ldots, d$; $i = 1, \ldots, 3$.

To implement the correct paired approach with unknown dispersion (as in BH-pair-correct and qvalue-pair-correct), we performed a two-sample negative binomial test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} \neq \mu_{Yj}$ using the coefficient from the negative binomial regression as the test statistic. Specifically, for each feature j we performed a negative binomial regression by treating the condition labels as a categorical covariate and feature j's measurements as the response. We implemented this regression analysis using function `glm.nb` in R package `MASS` and extracted the p-value of the coefficient as the p-value of feature j. The dispersion parameter was not pre-specified but estimated by `glm.nb`.

To implement the correct paired approach with known dispersion, we performed a similar negative binomial regression but with the pre-specified dispersion parameter 30^{-1} for each feature j. Then we computed the feature's p-value as the p-value of the coefficient of the condition covariate. We implemented this regression analysis using function `glm` in R package `stats`.

To implement the misspecified paired approach (as in BH-pair-mis and qvalue-pair-mis), we assumed that for each uninteresting feature j, $\{X_{ji}\}_{j=1}^3$ and $\{Y_{ji}\}_{j=1}^3$ follow the same Poisson distribution. We calculated the p-value of feature j from a two-sample Poisson test for the null hypothesis $H_0 : \mu_{Xj} = \mu_{Yj}$ against the alternative hypothesis $H_1 : \mu_{Xj} \neq \mu_{Yj}$ using function `poisson.test` in R package `stats`.

To implement the 2as1 paired approach (as in BH-pair-2as1 and qvalue-pair-2as1), we first used function `glm.nb` in R package `MASS` to estimate $\hat{\mu}_{Yj}$ and $\hat{\theta}_{Yj}$ from $\{Y_{ji}\}_{j=1}^3$. Then we computed the p-value of feature j by treating $\mathrm{NB}(3\hat{\mu}_{Bj}, (3\hat{\theta}_{Bj})^{-1})$ as the null distribution of $\sum_{i=1}^3 X_{ji}$ and calculated its two-sided tail probability, i.e., twice the smaller of the left-tail and right-tail probabilities.

S3 Bioinformatic methods with FDR control functionality

S3.1 Peak calling methods for ChIP-seq data

MACS2 MACS2 [1] uses sliding windows with a fixed length across the genome and identifies peaks by using a Poisson distribution to model the read counts within each window, which as one read count per replicate. Specifically, for each region (which is combined from sliding windows), MACS2 performs a one-sample Poisson test to calculate a p-value, where the null distribution is set to be Poisson with its parameter estimated from the background. By thresholding p-values, MACS2 identifies a set of candidate peaks. It also estimates for each candidate peak a q-value by swapping the experimental sample with the background (negative control) sample, and the q-values are used for FDR control. We used MACS2 software (version 2.2.6) with its default settings.

HOMER We used findPeaks, a program in HOMER [2], to perform peak calling on ChIP-seq data. The p-value calculation in findPeaks is similar to MACS2; that is, findPeaks also uses a Poisson distribution as the null distribution of read counts in a genomic region, and it also estimates the Poisson mean from the background sample. Then findPeaks identifies peaks by setting thresholds on p-values and fold-changes (the folder change of a region is defined as the observed read count under the experimental

sample divided by the estimated Poisson mean from the the background sample). We used findPeaks version 3.1.9.2.

S3.2 Mascot for peptide identification from MS data

Mascot Mascot uses probability-based scoring to identify PSMs from mass-spectrometry data. We ran Mascot in Proteome Discoverer 2.3.0.523 (ThermoScientific) with the following settings: 10 ppm precursor tolerance; 0.6 Da fragment tolerance; static modifications: methylthio (C); dynamic modifications: deamination (NQ), oxidation (M). We ran Percolator [55] in conjunction with Mascot with the target/decoy selection mode set to separate. For Mascot, for a range of target FDR thresholds ($q \in \{1\%, 2\%, \ldots, 10\%\}$), we identified the target PSMs with Mascot q-values no greater than q as discoveries. To prepare the input for Clipper, we set peptide and protein FDRs to 100% to obtain the entire lists of target PSMs and decoy PSMs with their Mascot q-values.

S3.3 Differentially expressed gene (DEG) methods for bulk RNA-seq data

edgeR edgeR models each gene's read counts by using a negative binomial regression, where the condition is incorporated as an indicator covariate, and the condition's coefficient represents the gene-wise differential expression effect [4]. We used `R` package `edgeR` version 3.30.0.

DESeq2 DESeq2 uses a similar negative binomial regression as edgeR to model each gene's read counts under two conditions. DESeq2 differs from edgeR mainly in their estimation of the dispersion parameter in the negative binomial distribution [5]. We used the `R` package `DESeq2` version 1.28.1.

S3.4 Differentially interacting chromatin regions (DIR) methods for Hi-C data

MultiHiCcompare MultiHiCcompare relies on a non-parametric method to jointly normalize multiple Hi-C interaction matrices [11]. It uses a generalized linear model to detect DIRs. MultiHiCcompare is an extension of the `HiCcompare` package [56]. We used `R` package `multiHiCcompare` version 1.6.0.

diffHic diffHic uses the statistical framework of the `edgeR` package to model biological variability and to test for significant differences between conditions [13]. We used `R` package `diffHic` version 1.20.0.

FIND FIND uses a spatial Poisson process to detect chromosomal regions that display a significant difference between two regions' interaction intensity and their neighbouring interaction intensities [12]. We used `R` package `FIND` version 0.99.

S4 Benchmark data generation in omics data applications

S4.1 ChIP-seq data with synthetic spike-in peaks

We used one experimental sample and two control samples (which we refer to as Control 1 and Control 2) from H3K4me3 ChIP-seq data in Chromosome 1 of the cell line GM12878 [45] .

(i) We input the experimental sample and Control 1 into MACS2 [1] to call peaks, which we call "tentative peaks."

(ii) Among tentative peaks, we further identified 25 "true peaks" as follows. We validated these peaks by comparing their read coverages in the experimental sample and Control 1 (Fig. S7).

(iii) We created a synthetic experimental sample by adding synthetic true peaks to Control 1. To mimic real H3K4me3 ChIP-seq data, where peaks are located predominantly in promoter regions, we added synthetic true peaks to promoter regions only. Specifically, we randomly sampled 585 genes' promoter regions from Chromosome 1. For each promoter region, we selected one of the 25 true peaks whose length was the closet to the length of the promoter region, and we down-sampled this peak's reads with probability p (that is, every read is sampled with probability p), which was drawn uniformly from 20% to 60%. Then we added the resulting down-sampled peak (i.e., its reads) as a synthetic true peak to that promoter region. After such addition for all the 587 promoter regions, we obtained a synthetic experimental sample.

(iv) We repeated Step (iii) for 20 times to generate 20 synthetic experimental samples. For each experimental sample, we paired it with Control 2 as the background sample. Hence, we obtained 20 synthetic ChIP-seq datasets, each containing 585 synthetic true peaks.

(v) After applying a peak calling method to these 20 synthetic datasets, we evaluated the 20 FDPs and 20 empirical power, which are then averaged as the approximate FDR and power of that peaking call method. In the evaluation, a called peak is a true positive if it overlapped with a synthetic true peak; otherwise, it is a false positive.

S4.2 Real MS standard data

The data generation information will be published in a future manuscript. Interested readers should contact Dr. Leo Wang at lewang@coh.org.

S4.3 Bulk RNA-seq data with synthetic spike-in DEGs

We used the human monocyte RNA-Seq dataset including 17 samples of classical monocytes and 17 samples of nonclassical monocytes [50]. Each sample contains expression levels of $d = 52,376$ genes.

(i) We first performed normalization on all 34 samples using the edgeR normalization method Trimmed Mean of M-values (TMM) [57]. We denote the resulting normalized read count matrix of classical and non-classical monocytes by $\mathbf{X}^{\mathrm{cl}}$ and $\mathbf{X}^{\mathrm{ncl}}$, respectively. Following the convention in bioinformatics, the columns and rows of $\mathbf{X}^{\mathrm{cl}}$ and $\mathbf{X}^{\mathrm{ncl}}$ represent biological samples and genes, respectively.

(ii) To define true DEGs, we first computed the log fold change of gene j by logfc$_j = \log_2\left[(\bar{\mathbf{X}}_j^{\mathrm{ncl}}+1)/(\bar{\mathbf{X}}_j^{\mathrm{cl}}+1)\right]$ for $j = 1,\ldots,d$, where $\mathbf{X}_i^{\mathrm{cl}}$ and $\mathbf{X}_i^{\mathrm{ncl}}$ denote the i-th row vector of $\mathbf{X}^{\mathrm{cl}}$ and $\mathbf{X}^{\mathrm{ncl}}$ respectively and $\bar{\cdot}$ denotes the average of elements in a vector. We added the pseudo-count of 1 to avoid division by 0. We defined true DEGs as those with logfc$_j \geq 4$ or logfc$_j \leq -4$.

(iii) We generated synthetic data with 3 samples under both the experimental and background conditions, a typical design in bulk RNA-seq experiments. Specifically, if gene j is a DEG, we randomly sampled without replacement 3 values from $\mathbf{X}_j^{\mathrm{cl}}$ as counts under the experimental condition, and another 3 values from $\mathbf{X}_j^{\mathrm{ncl}}$ as counts under the background condition. If gene j is not a DEG, we randomly sampled 6 values without replacement from $(\mathbf{X}_j^{\mathrm{cl}}, \mathbf{X}_j^{\mathrm{ncl}})$ and randomly split them into 3 and 3 counts under two conditions.

(iv) We repeated Step (iii) for 200 times to generate 200 synthetic datasets.

The human monocyte RNA-Seq data set is available in the NCBI Sequence Read Archive (SRA) under accession number SRP082682 (https://www.ncbi.nlm.nih.gov/Traces/study/?acc=srp082682).

S4.4 Hi-C data with synthetic spike-in DIRs

The real Hi-C interaction matrix contains the pairwise interaction intensities of 250 binned genomic regions in Chromosome 1. It is from the cell line GM12878 and available in the NCBI Gene Expression Omnibus(GEO) under accession number GSE63525. We denote the real interaction matrix as $\mathbf{X}^{\text{real}}$. Because $\mathbf{X}^{\text{real}}$ is symmetric, we only focus on its upper triangular part.

(i) Among the $(250 \times 250 - 250)/2 = 31,125$ upper triangular entries (i.e., region pairs), we selected 404 entries as true up-regulated DIRs, and 550 entries as true down-regulated DIRs (Fig. S8).

(ii) Next, for the (i, j)-th entry, we generated a log fold change, denoted by f_{ij}, between the two condition as follows. We simulated f_{ij} from truncated Normal$(100/|i - j|, 0.5^2)$ with support $[0.05, \infty)$ if the (i, j)-th entry is up-regulated, or from truncated Normal$(-100/|i - j|, 0.5^2)$ with support $(-\infty, -0.05]$ if the (i, j)-th entry is down-regulated; if the (i, j)-th entry is not differential, we set $f_{ij} = 0$.

(iii) Then we specify the mean measurement of the (i, j)-th entry under the two conditions as $\mu_{Xij} = [\mathbf{X}^{\text{real}}]_{ij}$ and $\mu_{Yij} = [\mathbf{X}^{\text{real}}]_{ij} \cdot e^{f_{ij}}$, respectively.

(iv) We generated synthetic read counts of the (i, j)-th entry from $\text{NB}(\mu_{Xij}, 1000^{-1})$ and $\text{NB}(\mu_{Yij}, 1000^{-1})$ respectively under the two conditions.

(v) We repeated Step (iv) for 200 times to generate 200 synthetic datasets.

S5 Implementation of Clipper in omics data applications

Below we briefly introduce the implementation of Clipper in the four omics data applications. All the results were obtained by running using `R` package `Clipper` (see package vignette for details: https://github.com/JSB-UCLA/Clipper/blob/master/vignettes/Clipper.pdf).

S5.1 Peak calling from ChIP-seq data

(i) We consider each genomic location, i.e., a base pair, as a feature and each ChIP-seq sample as a replicate under the experimental or background condition. Then we consider the read count of each location in each sample as the corresponding feature's measurement. Doing so, we summarized ChIP-seq data into a $d \times (m + n)$ matrix, where d is the number of locations, and m and n are the numbers of experimental and control samples, respectively. We then applied Clipper to perform an enrichment analysis to obtain the contrast score C_j for each location j. In our results, $m = n = 1$, so the default Clipper implementation is Clipper-diff-BC.

(ii) For any target FDR threshold q, Clipper gives a cutoff on contrast scores T_q.

(iii) We then used existing peak calling methods, e.g., MACS2 and HOMER, to call candidate peaks with the least stringent q-value cutoff. For example, when we used MACS2, we set the q-value cutoff as 1.

(iv) We computed the contrast score of each candidate peak as the median of the contrast scores of all the locations within.

(v) The candidate peaks with contrast scores greater than or equal to T_q are called discoveries.

S5.2 Peptide identification from mass spectrometry data

(i) We consider each mass spectrum as a feature and its target/decoy PSM as a replicate under the experimental/background condition respectively. Then we consider $-\log_{10}(\text{q-value} + 0.01)$ as the measurement of each PSM, where the q-value is output by Mascot. Doing so, we summarized the Mascot output into a $d \times (m+n)$ matrix, where d is the number of mass spectra, and m and n are the numbers of experimental and control samples, respectively. We then applied Clipper to perform an enrichment analysis to obtain a contrast score C_j for each mass spectrum j. If the mass spectrum has no decoy or background measurement, we set $C_j = 0$. In our results, $m = n = 1$, so the default Clipper implementation is Clipper-diff-BC.

(ii) For any target FDR threshold q, Clipper gives a cutoff on contrast scores T_q.

(iii) The target PSMs whose mass spectra have contrast scores greater than or equal to T_q are called discoveries.

S5.3 DEG identification from bulk RNA-seq data

(i) We consider each gene as a feature and the class label—classical and non-classical human monocytes—as the two conditions. Then we consider $\log_2$-transformed read counts with a pseudocount 1 as measurements. Doing so, we summarized the gene expression matrix into a $d \times (m+n)$ matrix, where d is the number of genes, and m and n are the numbers of samples under the two conditions, respectively. We then applied Clipper to perform a differential analysis to obtain a contrast score C_j for each gene. In our results, $m = n = 3$, so the default Clipper implementation is Clipper-max-GZ with $h = 1$.

(ii) For any target FDR threshold q, Clipper gives a cutoff on contrast scores T_q.

(iii) The genes with contrast scores greater than or equal to T_q are called discoveries.

S5.4 DIR identification from Hi-C data

(i) We consider each pair of genomic regions as a feature and manually created two conditions. Then we consider $\log$-transformed read counts as measurements. Doing so, we summarized the gene expression matrix into a $d \times (m+n)$ matrix, where d is the total pairs of genomic regions, and m and n are the numbers of samples under the two conditions, respectively. We then applied Clipper to perform a differential analysis to obtain a contrast score C_j for each pair of genomic regions. In our results, $m = n = 2$, so the default Clipper implementation is Clipper-max-GZ with $h = 1$.

(ii) For any target FDR threshold q, Clipper gives a cutoff on contrast scores T_q.

(iii) The pairs of genomic regions with contrast scores greater than or equal to T_q are called discoveries.

S6 Proofs

S6.1 Proof of Theorem 1

We first prove Theorem 1, which relies on Lemmas 1 and 2. Here we only include the proof of Lemma 1 and defer the proof of Lemma 2 to Section S6.3.

Proof 1 (Proof of Lemma 1) *Here we prove that Lemma 1 holds when C_j is constructed using (5); the proof is similar when C_j is constructed using (6).*

When input data satisfy (1) and (2) and $m = n$, properties (a) and (b) can be derived directly. To prove property (c), it suffices to prove that for any $j \in \mathcal{N}$ with $C_j \neq 0$, S_j is independent of $|C_j|$.

Note that $\bar{X}_j$ and $\bar{Y}_j$ are i.i.d for $j \in \mathcal{N}$ when $m = n$. Hence for any measurable set $\mathcal{A} \subset [0, +\infty)$,

$$\mathbb{P}\left(S_j = 1, |C_j| \in \mathcal{A}\right) = \mathbb{P}\left(t^{\text{diff}}(\boldsymbol{X}_j, \boldsymbol{Y}_j) \in \mathcal{A}\right) = \mathbb{P}\left(t^{\text{diff}}(\boldsymbol{Y}_j, \boldsymbol{X}_j) \in \mathcal{A}\right) = \mathbb{P}\left(S_j = -1, |C_j| \in \mathcal{A}\right).$$

The first equality holds because $t^{\text{diff}}(\boldsymbol{X}_j, \boldsymbol{Y}_j) = C_j = |C_j|$ when $S_j = 1$. The second equality holds because $t^{\text{diff}}(\boldsymbol{X}_j, \boldsymbol{Y}_j)$ and $t^{\text{diff}}(\boldsymbol{Y}_j, \boldsymbol{X}_j)$ are identically distributed when $j \in \mathcal{N}$. The third equality holds because $t^{\text{diff}}(\boldsymbol{Y}_j, \boldsymbol{X}_j) = -C_j$; if $-C_j \in \mathcal{A}$, then $S_j = -1$.

Because $\mathbb{P}\left(S_j = 1, |C_j| \in \mathcal{A}\right) + \mathbb{P}\left(S_j = -1, |C_j| \in \mathcal{A}\right) = \mathbb{P}\left(|C_j| \in \mathcal{A}\right)$, it follows that

$$\mathbb{P}\left(S_j = 1, |C_j| \in \mathcal{A}\right) = \frac{1}{2}\mathbb{P}\left(|C_j| \in \mathcal{A}\right) = \mathbb{P}(S_j = 1)\mathbb{P}\left(|C_j| \in \mathcal{A}\right),$$

where the last equality holds because $\mathbb{P}(S_j = 1) = 1/2$ by property (b).

Hence, S_j and $|C_j|$ are independent $\forall j \in \mathcal{N}$.

Proof 2 (Proof of Theorem 1) *Define a random subset of $\mathcal{N}$ as $\mathcal{M} := \mathcal{N} \backslash \{j \in \mathcal{N} : C_j = 0\} = \{j \in \mathcal{N} : S_j \neq 0\}$.*

First note that by Lemma 1(b), $\mathbb{P}(S_j = -1) = \mathbb{P}(C_j < 0) = 1/2$ for all $j \in \mathcal{M} \subset \mathcal{N}$. Assume without loss of generality that $\mathcal{M} = \{1, \ldots, d'\}$. We order $\{|C_j| : j \in \mathcal{M}\}$, from the largest to the smallest, denoted by $|C_{(1)}| \geq |C_{(2)}| \geq \cdots \geq |C_{(d')}|$. Let $J = \sum_{j \in \mathcal{N}} \mathbb{1}\left(|C_j| \geq T^{\text{BC}}\right)$, the number of uninteresting features whose contrast scores have absolute values no less than T^{BC}. When $J > 0$, $|C_{(1)}| \geq \cdots \geq |C_{(J)}| \geq T^{\text{BC}}$. Define $Z_k = \mathbb{1}\left(C_{(k)} < 0\right)$, $k = 1, \ldots, d'$. Then for each order k, the following holds

$$C_{(k)} \geq T^{\text{BC}} \Longleftrightarrow \left|C_{(k)}\right| \geq T^{\text{BC}} \text{ and } C_{(k)} > 0 \Longleftrightarrow k \leq J \text{ and } Z_k = 0\,;$$
$$C_{(k)} \leq -T^{\text{BC}} \Longleftrightarrow \left|C_{(k)}\right| \geq T^{\text{BC}} \text{ and } C_{(k)} < 0 \Longleftrightarrow k \leq J \text{ and } Z_k = 1\,.$$

Then

$$\begin{aligned}
\frac{\text{card}(\{j \in \mathcal{M} : C_j \geq T^{\text{BC}}\})}{\text{card}(\{j \in \mathcal{M} : C_j \leq -T^{\text{BC}}\}) + 1} &= \frac{\sum_{k=1}^{d'} \mathbb{1}\left(C_{(k)} \geq T^{\text{BC}}\right)}{1 + \sum_{k=1}^{d'} \mathbb{1}\left(C_{(k)} \leq -T^{\text{BC}}\right)} \\
&= \frac{\sum_{k=1}^{J} \mathbb{1}\left(C_{(k)} \geq T^{\text{BC}}\right)}{1 + \sum_{k=1}^{J} \mathbb{1}\left(C_{(k)} \leq -T^{\text{BC}}\right)} \\
&= \frac{(1 - Z_1) + \cdots + (1 - Z_J)}{1 + Z_1 + \cdots + Z_J} \\
&= \frac{1 + J}{1 + Z_1 + \cdots + Z_J} - 1\,.
\end{aligned}$$

Because $\{S_j\}_{j \in \mathcal{N}}$ is independent of $\mathcal{C}$ (Lemma 1(c)), Lemma 1(a)-(b) still holds after $C_1, \ldots, C_{d'}$ are

reordered as $C_{(1)}, \ldots, C_{(d')}$. *Thus* $Z_1, \ldots, Z_{d'}$ *are i.i.d. from* $\mathrm{Bernoulli}(1/2)$. *To summarize, it holds that*

$$\{Z_j\}_{j \in \mathcal{M}} \,\Big|\, \mathcal{M} \overset{\text{i.i.d.}}{\sim} \mathrm{Bernoulli}(1/2)\,.$$

Then by applying Lemma 2 and making $\rho = 0.5$, *we have:*

$$\mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{M} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j \in \mathcal{M} : C_j \leq -T^{\mathrm{BC}}\}) + 1} \,\Bigg|\, \mathcal{M}\right] \leq 1 \tag{16}$$

Then

$$\begin{aligned}
\mathrm{FDR} &= \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{BC}}\}) \vee 1}\right] \\
&= \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}) + 1} \cdot \frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}) + 1}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{BC}}\}) \vee 1}\right] \\
&\leq \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}) + 1} \cdot \frac{\mathrm{card}(\{j : C_j \leq -T^{\mathrm{BC}}\}) + 1}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{BC}}\}) \vee 1}\right] \\
&\leq q \cdot \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{BC}}\}) + 1}\right] \\
&\leq q \cdot \mathbb{E}\left[\mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{M} : C_j \geq T^{\mathrm{BC}}\})}{\mathrm{card}(\{j \in \mathcal{M} : C_j \leq -T^{\mathrm{BC}}\}) + 1} \,\Bigg|\, \mathcal{M}\right]\right] \\
&\leq q,
\end{aligned}$$

where $\mathcal{M}$ *is random subset of* $\mathcal{N}$ *such that for each* $j \in \mathcal{M}$, $|C_j| > 0$. *The last inequality follows from* (16).

S6.2 Proof of Theorem 2

We then prove Theorem 2, which relies on Lemmas 2 and 3. Here we introduce the proof of Lemma 3 and defer the proof of Lemma 2 to Section S6.3.

Proof 3 (Proof of Lemma 3) *With input data satisfying* (1) *and* (2), C_j *constructed from* (10) *or* (11), *property (a) can be derived directly.*

To show property (b), note that for each uninteresting feature $j \in \mathcal{N}$, $\boldsymbol{X}_j$ *and* $\boldsymbol{Y}_j$ *are from the same distribution; thus* $\{T_j^{\sigma_\ell}\}_{\ell=0}^h$ *are identically distributed. Define an event* $\mathcal{E}_j := \left\{\sum_{\ell=0}^h \mathbb{1}(T_j^{\sigma_\ell} = T_j^{(0)}) = 1\right\}$, *which indicates that* $T_j^{(0)}$, *the maximizer of* $\{T_j^{\sigma_\ell}\}_{\ell=0}^h$, *is unique. Then conditional on* $\mathcal{E}_j$, *the maximizer is equally likely to be any of* $\{0, \ldots, h\}$, *and it follows that* $\mathbb{P}(S_j = 1 \mid \mathcal{E}_j) = \mathbb{P}(T_j^{\sigma_0} = T_j^{(0)} \mid \mathcal{E}_j) = 1/(h+1)$. *Conditioning on that* $\mathcal{E}_j$ *does not happen,* $\mathbb{P}(S_j = 1 \mid \mathcal{E}_j^c) = 0$. *Thus* $\mathbb{P}(S_j = 1) = \mathbb{P}(S_j = 1 \mid \mathcal{E}_j)\mathbb{P}(\mathcal{E}_j) + \mathbb{P}(S_j = 1 \mid \mathcal{E}_j^c)\mathbb{P}(\mathcal{E}_j^c) \leq 1/(h+1)$.

The proof of property (c) is similar to the Proof of Lemma 1(c). It suffices to show that for any $j \in \mathcal{N}$ *with* $C_j \neq 0$ *(that is,* $\mathcal{E}_j$ *occurs),* S_j *is independent of* $|C_j|$. *As* $\boldsymbol{X}_j$ *and* $\boldsymbol{Y}_j$ *are from the same distribution,* $\{T_j^{\sigma_\ell}\}_{\ell=0}^h$ *are identically distributed. Hence for any measurable set* $\mathcal{A} \subset [0, +\infty)$,

$$\begin{aligned}
\mathbb{P}\left(S_j = 1, |C_j| \in \mathcal{A} \mid \mathcal{E}_j\right) &= \mathbb{P}\left(T_j^{\sigma_0} = T_j^{(0)}, |C_j| \in \mathcal{A} \,\middle|\, \mathcal{E}_j\right) \\
&= \frac{1}{h}\mathbb{P}\left(T_j^{\sigma_0} \neq T_j^{(0)}, |C_j| \in \mathcal{A} \,\middle|\, \mathcal{E}_j\right) \\
&= \frac{1}{h}\mathbb{P}\left(S_j = -1, |C_j| \in \mathcal{A} \mid \mathcal{E}_j\right).
\end{aligned}$$

The first equality holds because $T_j^{\sigma_0} = T_j^{(0)}$ when $S_j = 1$. The second equality holds because $\{T_j^{\sigma_\ell}\}_{\ell=0}^h$ are identically distributed when $j \in \mathcal{N}$. The third equality holds because $T_j^{\sigma_0} \neq T_j^{(0)}$ when $S_j = -1$.

Because $\mathbb{P}\left(S_j = 1, |C_j| \in \mathcal{A} \mid \mathcal{E}_j\right) + \mathbb{P}\left(S_j = -1, |C_j| \in \mathcal{A} \mid \mathcal{E}_j\right) = \mathbb{P}\left(|C_j| \in \mathcal{A} \mid \mathcal{E}_j\right)$, it follows that

$$\mathbb{P}\left(S_j = 1, |C_j| \in \mathcal{A} \mid \mathcal{E}_j\right) = \frac{1}{h+1}\mathbb{P}\left(|C_j| \in \mathcal{A} \mid \mathcal{E}_j\right) = \mathbb{P}(S_j = 1 \mid \mathcal{E}_j)\mathbb{P}\left(|C_j| \in \mathcal{A} \mid \mathcal{E}_j\right) ,$$

where the last equality holds because $\mathbb{P}(S_j = 1 \mid \mathcal{E}_j) = 1/(h+1)$.

Hence, S_j and $|C_j|$ are independent $\forall j \in \mathcal{N}$ with $C_j \neq 0$.

Proof 4 (Proof of Theorem 2) *Define a random subset of $\mathcal{N}$ as $\mathcal{M} := \mathcal{N} \backslash \{j \in \mathcal{N} : C_j = 0\} = \{j \in \mathcal{N} : S_j \neq 0\}$. Assume without loss of generality that $\mathcal{M} = \{1, \ldots, d'\}$. We order $\{|C_j| : j \in \mathcal{M}\}$, from the largest to the smallest, denoted by $|C_{(1)}| \geq |C_{(2)}| \geq \cdots \geq |C_{(d')}|$. Let $J = \sum_{j \in \mathcal{N}} \mathbb{1}\left(|C_j| \geq T^{\mathrm{GZ}}\right)$, the number of uninteresting features whose contrast scores have absolute values no less than T^{GZ}. When $J > 0$, $|C_{(1)}| \geq \cdots \geq |C_{(J)}| \geq T^{\mathrm{GZ}}$. Define $Z_k = \mathbb{1}\left(C_{(k)} < 0\right)$, $k = 1, \ldots, d'$. Then for each order k, the following holds:*

$$C_{(k)} \geq T^{\mathrm{GZ}} \Longleftrightarrow \left|C_{(k)}\right| \geq T^{\mathrm{GZ}} \text{ and } C_{(k)} > 0 \Longleftrightarrow k \leq J \text{ and } Z_k = 0\,;$$
$$C_{(k)} \leq -T^{\mathrm{GZ}} \Longleftrightarrow \left|C_{(k)}\right| \geq T^{\mathrm{GZ}} \text{ and } C_{(k)} < 0 \Longleftrightarrow k \leq J \text{ and } Z_k = 1\,.$$

Then it follows that

$$\begin{aligned}
\frac{\mathrm{card}(\{j \in \mathcal{M} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j \in \mathcal{M} : C_j \leq -T^{\mathrm{GZ}}\}) + 1} &= \frac{\sum_{k=1}^{d'} \mathbb{1}\left(C_{(k)} \geq T^{\mathrm{GZ}}\right)}{1 + \sum_{k=1}^{d'} \mathbb{1}\left(C_{(k)} \leq -T^{\mathrm{GZ}}\right)} \\
&= \frac{\sum_{k=1}^{J} \mathbb{1}\left(C_{(k)} \geq T^{\mathrm{GZ}}\right)}{1 + \sum_{k=1}^{J} \mathbb{1}\left(C_{(k)} \leq -T^{\mathrm{GZ}}\right)} \\
&= \frac{(1 - Z_1) + \cdots + (1 - Z_J)}{1 + Z_1 + \cdots + Z_J} \\
&= \frac{1 + J}{1 + Z_1 + \cdots + Z_J} - 1\,.
\end{aligned}$$

Because $\{S_j\}_{j \in \mathcal{N}}$ is independent of $\mathcal{C}$ (Lemma 1(c)), Lemma 1(a)-(b) still holds after $C_1, \ldots, C_{d'}$ are reordered as $C_{(1)}, \ldots, C_{(d')}$. Thus $Z_1, \ldots, Z_{d'}$ are i.i.d. from $\mathrm{Bernoulli}(\rho_k)$. To summarize, it holds that

$$\{Z_j\}_{j \in \mathcal{M}} \,\Big|\, \mathcal{M} \overset{\text{i.i.d.}}{\sim} \mathrm{Bernoulli}(\rho_k)\,.$$

Then by applying Lemma 2 and making $\rho = h/(h+1)$, we have:

$$\mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{M} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j \in \mathcal{M} : C_j \leq -T^{\mathrm{GZ}}\}) + 1}\right] \leq 1/h\,. \tag{17}$$

Then

$$
\begin{aligned}
\mathrm{FDR} &= \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{GZ}}\}) \vee 1}\right] \\
&= \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{GZ}}\}) + 1} \cdot \frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{GZ}}\}) + 1}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{GZ}}\}) \vee 1}\right] \\
&\leq h \cdot \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{GZ}}\}) + 1} \cdot \frac{\frac{1}{h}\mathrm{card}(\{j : C_j \leq -T^{\mathrm{GZ}}\}) + \frac{1}{h}}{\mathrm{card}(\{j : C_j \geq T^{\mathrm{GZ}}\}) \vee 1}\right] \\
&\leq hq \cdot \mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{N} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j \in \mathcal{N} : C_j \leq -T^{\mathrm{GZ}}\}) + 1}\right] \\
&\leq hq \cdot \mathbb{E}\left[\mathbb{E}\left[\frac{\mathrm{card}(\{j \in \mathcal{M} : C_j \geq T^{\mathrm{GZ}}\})}{\mathrm{card}(\{j \in \mathcal{M} : C_j \leq -T^{\mathrm{GZ}}\}) + 1} \,\middle|\, \mathcal{M}\right]\right] \\
&\leq q\,,
\end{aligned}
$$

where the second inequality follows from the definition of T_{GZ} (12) *and the last inequality follows from* (17)*.*

S6.3 Proof of Lemma 2

Finally, we derive Lemma 2 by following the same proof same as in [37], which relies on Lemma 4 and Corollary 1.

Lemma 4 *Suppose that* $Z_1, \ldots, Z_d \overset{i.i.d.}{\sim} \mathrm{Bernoulli}(\rho)$*. Let* J *be a stopping time in reverse time with respect to the filtration* $\{\mathcal{F}_j\}$*, where* $\mathcal{F}_j = \sigma\left(\{(Z_1 + \cdots + Z_j),\, Z_{j+1},\, \cdots,\, Z_d\}\right)$ *with* $\sigma(\cdot)$ *denoting a* σ*-algebra, and the variables* $Z_1, \ldots, Z_j$ *are exchangeable with respect to* $\{\mathcal{F}_j\}$*. Then*

$$
\mathbb{E}\left[\frac{1 + J}{1 + Z_1 + \cdots + Z_J}\right] \leq \rho^{-1}.
$$

Proof 5 (Proof of Lemma 4) *Define*

$$
Y_j = Z_1 + \cdots + Z_j \in \mathcal{F}_j
$$

and define the process

$$
M_j = \frac{1 + j}{1 + Z_1 + \cdots + Z_j} = \frac{1 + j}{1 + Y_j} \in \mathcal{F}_j.
$$

In [36], it is shown that $\mathbb{E}[M_d] \leq \rho^{-1}$*. Therefore, by the optional stopping time theorem it suffices to show that* $\{M_j\}$ *is a supermartingale with respect to* $\{\mathcal{F}_j\}$*. As* $\{Z_1, \ldots, Z_{j+1}\}$ *are exchangeable with respect to* $\mathcal{F}_{j+1}$*, we have*

$$
\mathbb{P}\left(Z_{j+1} = 1 \,|\, \mathcal{F}_{j+1}\right) = \frac{Y_j + 1}{1 + j}.
$$

Therefore, if $Y_{j+1} > 0$,

$$\begin{aligned}
\mathbb{E}\left[M_j \mid \mathcal{F}_{j+1}\right] &= \frac{1+j}{1+Y_{j+1}} \cdot \mathbb{P}\left(Z_{j+1}=0 \mid \mathcal{F}_{j+1}\right) + \frac{1+j}{1+Y_{j+1}-1} \cdot \mathbb{P}\left(Z_{j+1}=1 \mid \mathcal{F}_{j+1}\right) \\
&= \frac{1+j}{1+Y_{j+1}} \cdot \frac{1+j-Y_{j+1}}{1+j} + \frac{1+j}{1+Y_{j+1}-1} \cdot \frac{Y_{j+1}}{1+j} \\
&= \frac{1+j-Y_{j+1}}{1+Y_{j+1}} + 1 \\
&= \frac{1+(j+1)}{1+Y_{j+1}} \\
&= M_{j+1}.
\end{aligned}$$

If instead $Y_{j+1} = 0$, *then trivially* $Y_j = 0$, *and* $M_j = 1 + j < 2 + j = M_{j+1}$. *This proves that* $\{M_j\}$ *is a supermartingale with respect to* $\{\mathcal{F}_j\}$ *as desired.*

Corollary 1 *Suppose that* $\mathcal{A} \subseteq \{1, \ldots, d\}$ *is fixed, while* $Z_1, \ldots, Z_d \overset{i.i.d.}{\sim} \mathrm{Bernoulli}(\rho)$. *Let* J *be a stopping time in reverse time with respect to the filtration* $\{\mathcal{F}_j\}$, *where* $\mathcal{F}_j = \sigma\left(\{\sum_{k \le j, k \in \mathcal{A}} Z_k\} \cup \{Z_k : j < k < d, k \in \mathcal{A}\}\right)$ *with* $\sigma(\cdot)$ *denoting a* σ*-algebra, and the variables* $\{Z_k : k \le j, k \in \mathcal{A}\}$ *are exchangeable with respect to* $\mathcal{F}_j$. *Then*

$$\mathbb{E}\left[\frac{1+\mathrm{card}\left(\{k : k \le J, k \in \mathcal{A}\}\right)}{1+\sum_{k \le J, k \in \mathcal{A}} Z_k}\right] \le \rho^{-1}.$$

Proof 6 (Proof of Corollary 1) *Let* $\mathcal{A} = \{j_1, \ldots, j_m\}$ *where* $1 \le j_1 < \cdots < j_m \le d$. *Then by considering the i.i.d. sequence*

$$Z_{j_1}, \ldots, Z_{j_m}$$

in place of $Z_1, \ldots, Z_d$, *we see that this result is equivalent to Lemma 4.*

Proof 7 (Proof of Lemma 2) *[From [36]] We may assume* $\rho < 1$ *to avoid the trivial case. We first introduce a different definition for* $\{Z_j\}_{j=1}^d$ *by defining a random set* $\mathcal{A} \subseteq \{1, \ldots, d\}$ *where for each* j, *independently,*

$$\mathbb{P}\left(j \in \mathcal{A}\right) = \frac{1-\rho_j}{1-\rho}.$$

We then define random variables $Q_1, \ldots, Q_d \overset{i.i.d.}{\sim} \mathrm{Bernoulli}(\rho)$, *which are generated independently of the random set* $\mathcal{A}$. *Finally, we define*

$$Z_j = Q_j \cdot \mathbb{1}\left(j \in \mathcal{A}\right) + \mathbb{1}\left(j \notin \mathcal{A}\right). \tag{18}$$

Then $\{Z_j\}_{k=1}^d$ *are mutually independent and* $\mathbb{P}\left(Z_j = 1\right) = 1 - \mathbb{P}(j \in \mathcal{A}) \cdot \mathbb{P}(Q_j = 0) = \rho_j$, *that is,* $Z_j \sim \mathrm{Bernoulli}(\rho_j)$. *This new definition of* $\{Z_j\}_{j=1}^d$ *meet all the conditions required by Lemma 2, so that we can apply this new definition in the following proof.*

As $Z_j = Q_j \cdot \mathbb{1}\left(j \in \mathcal{A}\right) + \mathbb{1}\left(j \notin \mathcal{A}\right)$ *for all* j, *we have*

$$\frac{1+J}{1+Z_1+\cdots+Z_J} = \frac{1+\mathrm{card}\left(\{j \le J : j \in \mathcal{A}\}\right) + \mathrm{card}\left(\{j \le J : j \notin \mathcal{A}\}\right)}{1+\sum_{j \le J, j \in \mathcal{A}} Q_j + \mathrm{card}\left(\{j \le J : j \notin \mathcal{A}\}\right)} \le \frac{1+\mathrm{card}\left(\{j \le J : j \in \mathcal{A}\}\right)}{1+\sum_{j \le J, j \in \mathcal{A}} Q_j}, \tag{19}$$

where the last step uses the identify $\frac{a+c}{b+c} \le \frac{a}{b}$ *whenever* $0 < b \le a$ *and* $c \ge 0$. *Therefore, it will be sufficient to prove that*

$$\mathbb{E}\left[\frac{1+\mathrm{card}\left(\{j \le J : j \in \mathcal{A}\}\right)}{1+\sum_{j \le J, j \in \mathcal{A}} Q_j} \,\middle|\, \mathcal{A}\right] \le \rho^{-1}, \tag{20}$$

To prove (20), first let $\widetilde{Q}_j = Q_j \cdot \mathbb{1}(j \in \mathcal{A})$, and define a filtration $\{\mathcal{F}'_j\}$ where $\mathcal{F}'_j$ is the σ-algebra generated as

$$\mathcal{F}'_j = \sigma\left(\left\{\widetilde{Q}_1 + \cdots + \widetilde{Q}_j, \widetilde{Q}_{j+1}, \cdots, \widetilde{Q}_d, \mathcal{A}\right\}\right).$$

Next for any j, by (18) we see that

$$Z_1 + \cdots + Z_j, Z_{j+1}, \ldots, Z_d \in \mathcal{F}'_j \Rightarrow \mathcal{F}_j \subseteq \mathcal{F}'_j,$$

so J is a stopping time (in reverse time) with respect to $\mathcal{F}'_j$. Finally, since the Q_j's are independent of $\mathcal{A}$, (20) follows from Corollary 1 after conditioning on $\mathcal{A}$.

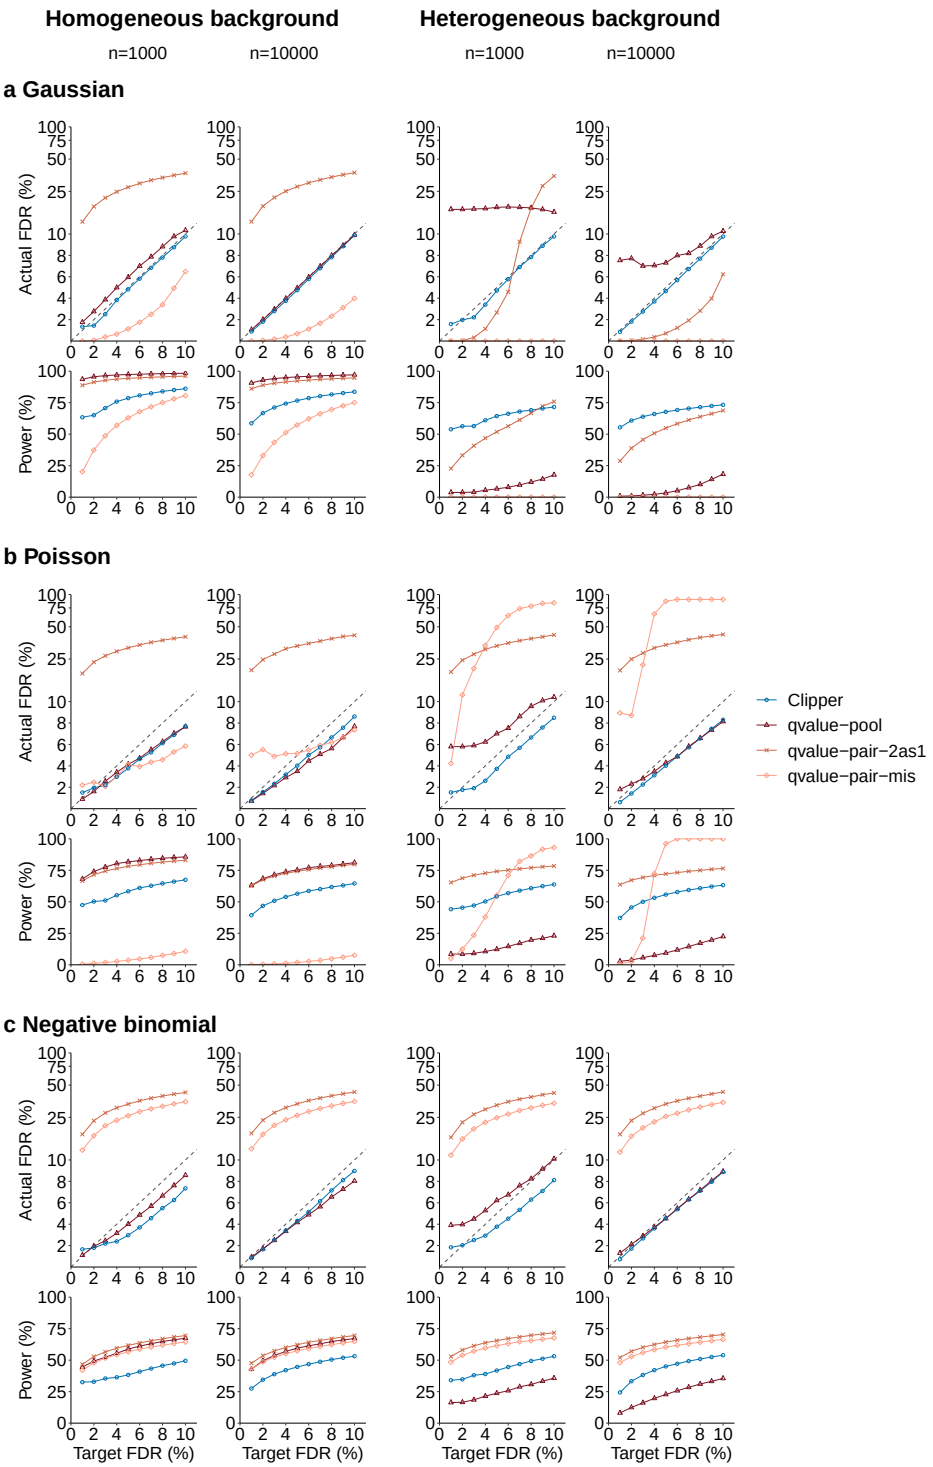

Figure S1: In 1vs1 enrichment analysis, comparison of Clipper and three other generic FDR control methods using Storey's q-value (qvalue-pool, qvalie-pair-2as1, and qvalue-pair-mis) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with $d = 1000$ or 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background.

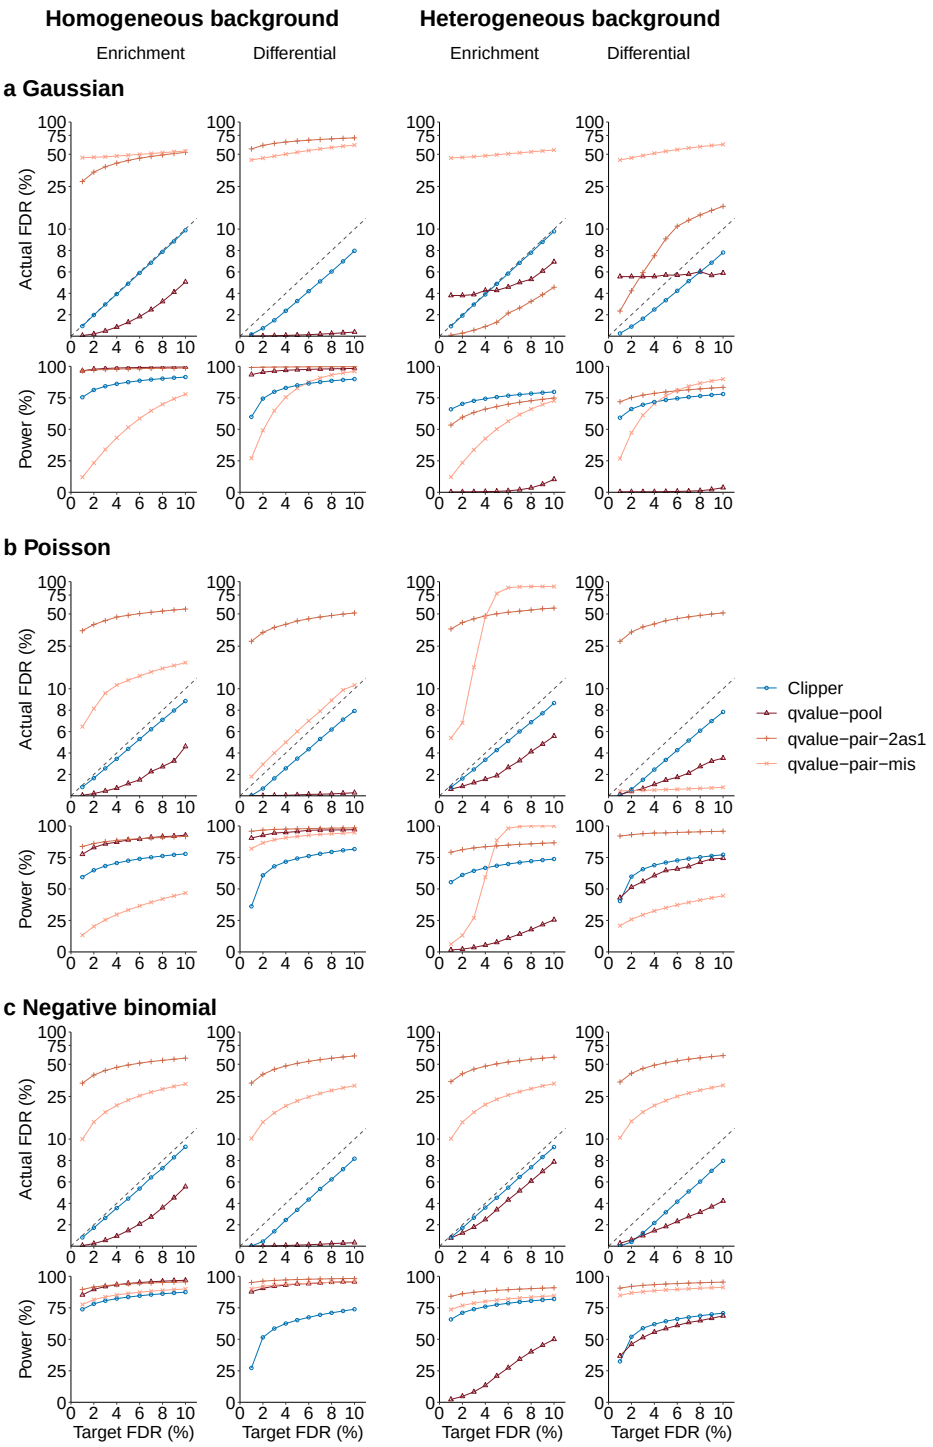

Figure S2: In 2vs1 enrichment analysis (columns 1 and 3) and differential analysis (columns 2 and 4), comparison of Clipper and three other generic FDR control methods using Storey's q-value (qvalue-pool, qvalie-pair-2as1, and qvalue-pair-mis) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background.

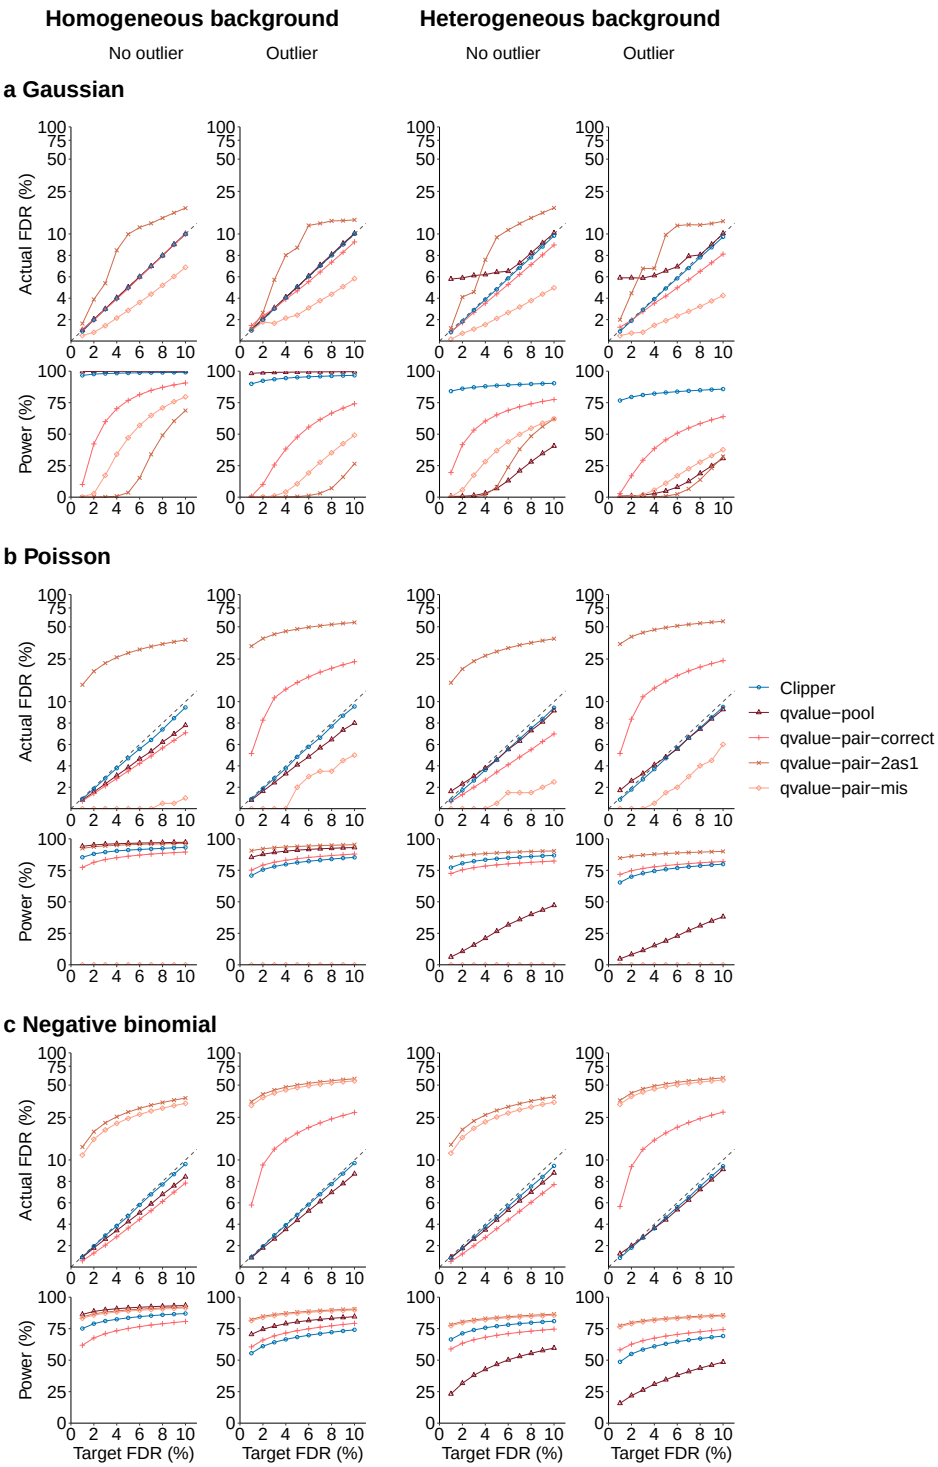

Figure S3: In 3vs3 enrichment analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of Clipper and four other generic FDR control methods using Storey's q-value (qvalue-pooled, qvalue-pair-correct, qvalue-pair-2as1, and qvalue-pair-mis) in terms of their FDR control and power in 3vs3 enrichment analysis with possible outliers. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background.

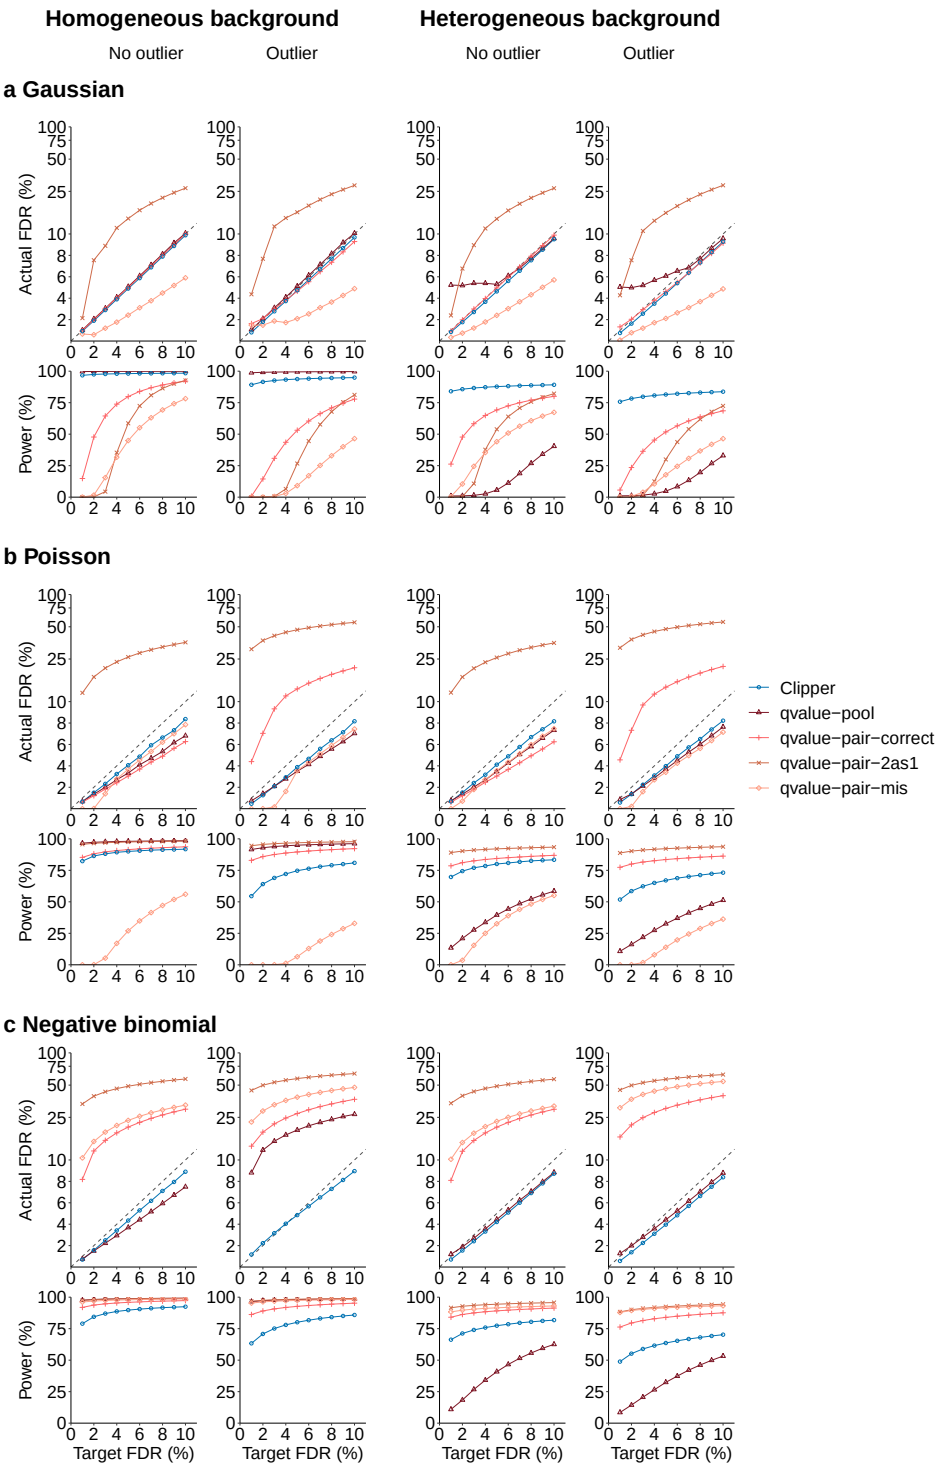

Figure S4: In 3vs3 differential analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of Clipper and four other generic FDR control methods using Storey's q-value (qvalue-pooled, qvalue-pair-correct, qvalue-pair-2as1, and qvalue-pair-mis) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Among the methods that control the FDR, Clipper is the second most powerful for homogeneous background and the most powerful for heterogeneous background (except for Poisson distribution where Clipper is second to qvalue-pair-correct, an idealistic method).

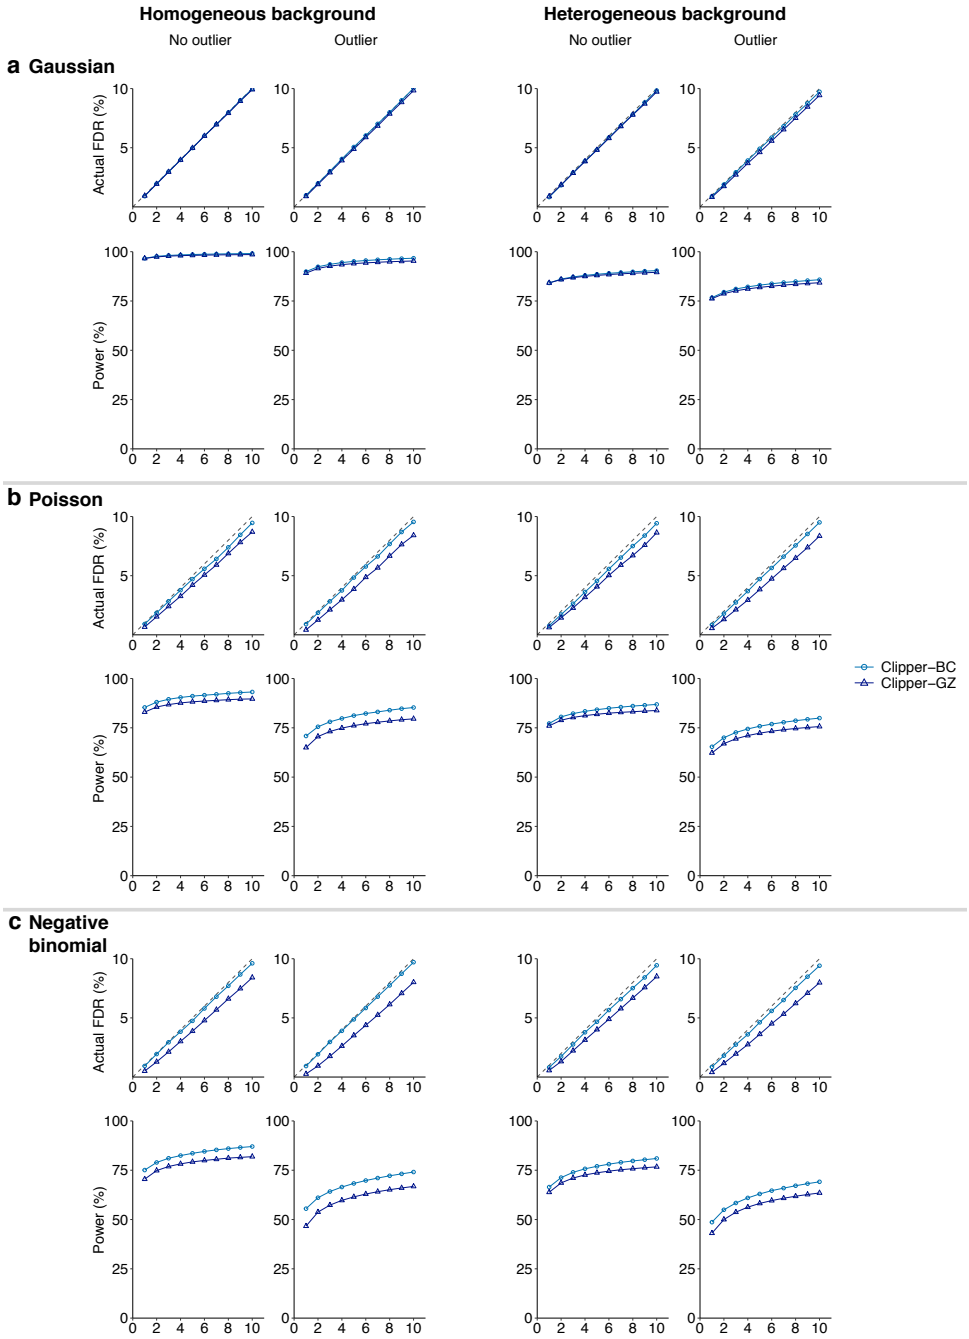

Figure S5: In 3vs3 enrichment analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of two Clipper variant algorithms (Clipper-diff-BC, Clipper-max-GZ(h=1)) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. Clipper-diff-BC is chosen as the default implementation under this scenario.

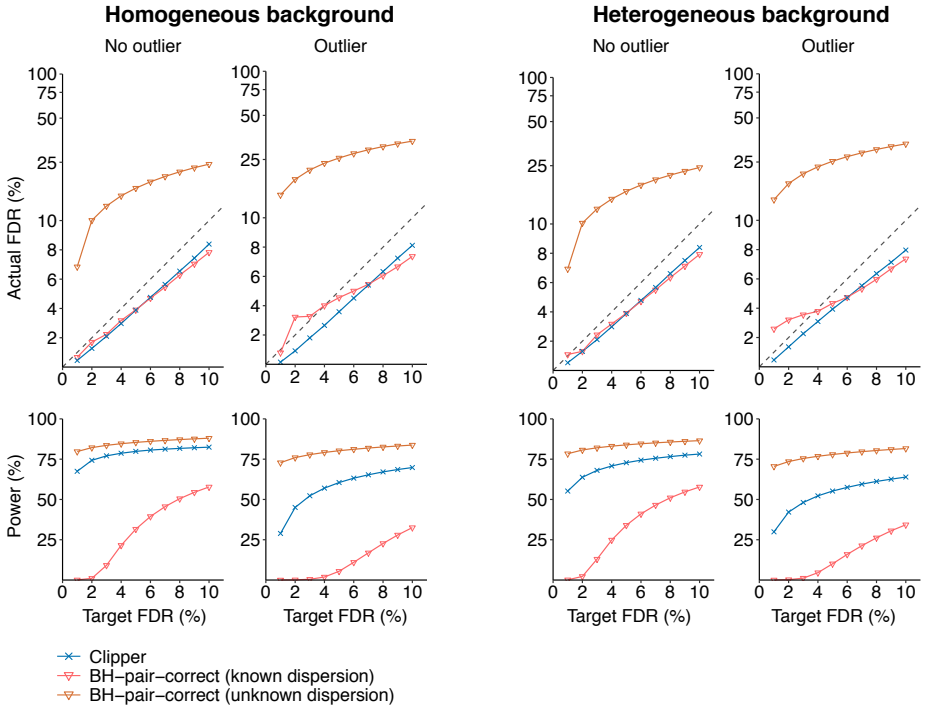

Figure S6: In 3vs3 differential analysis without (columns 1 and 3) or with outliers (columns 2 and 4), comparison of Clipper, BH-pair-correct (known dispersion), and BH-pair-correct (unknown dispersion) in terms of their FDR control and power. At target FDR thresholds $q \in \{1\%, 2\%, \cdots, 10\%\}$, each algorithm's actual FDRs and power are evaluated on 200 simulated datasets with 10000 features generated from (**a**) the Gaussian distribution, (**b**) the Poisson distribution, or (**c**) the negative binomial distribution under homogeneous (two left columns) and heterogeneous (two right columns) background scenarios. BH-pair-correct (unknown dispersion) cannot control the FDR in all settings. In contrast, Clipper is consistently the most powerful for homogeneous and heterogeneous background.

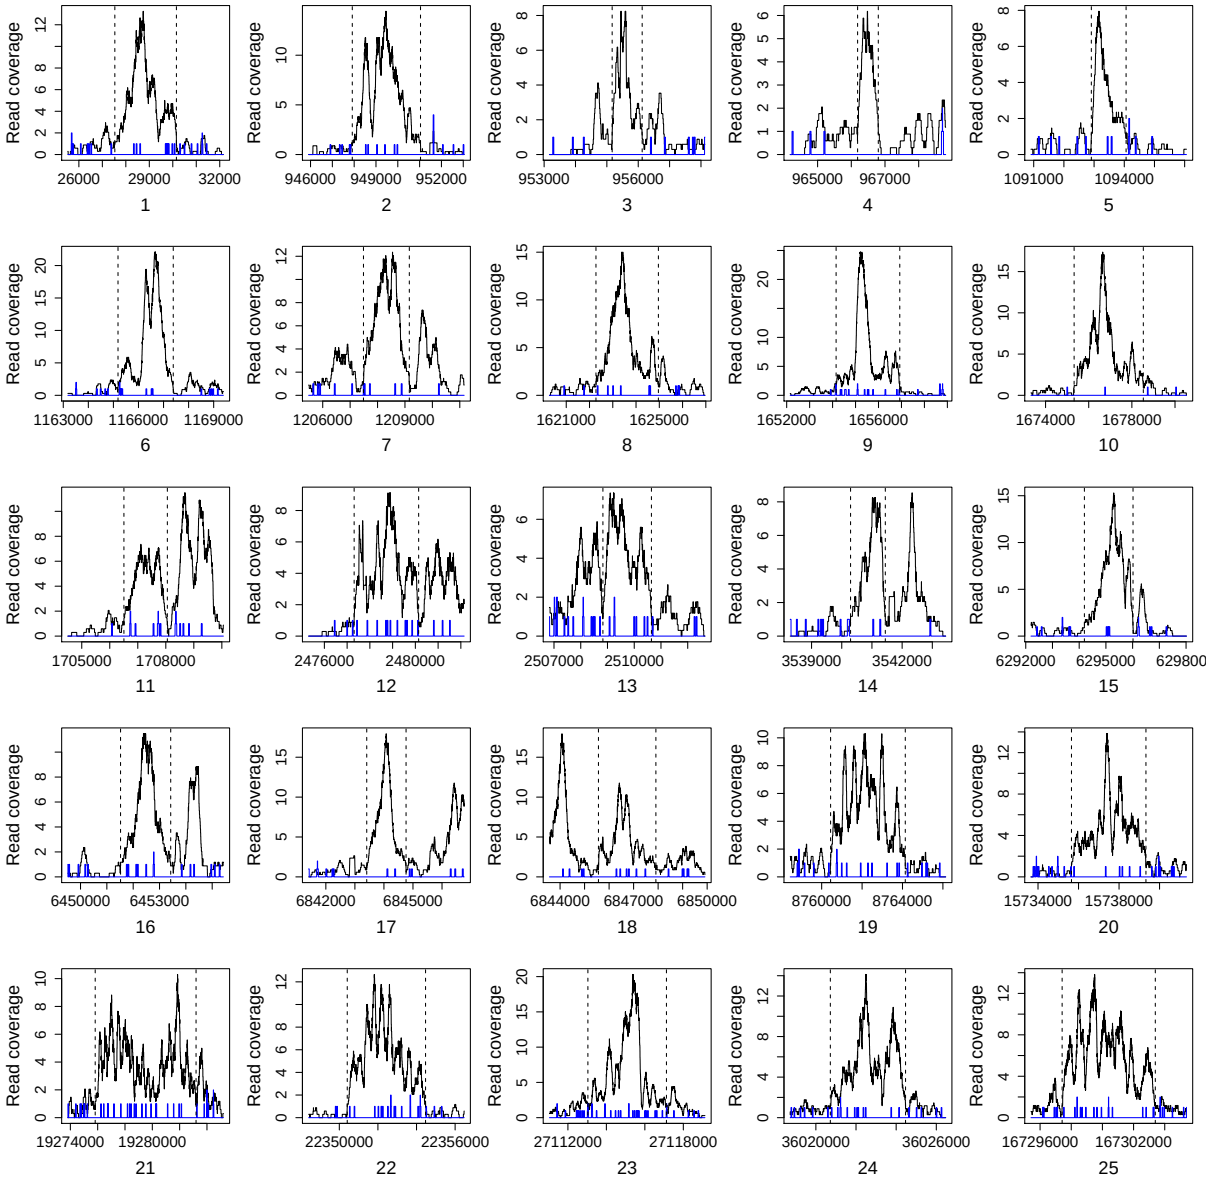

Figure S7: 25 "true peaks" from H3K4me3 ChIP-seq data of cell line GM12878. Black and blue curves indicate the read coverages in the experimental and control samples, respectively. Vertical dashed lines indicate the peak boundaries.

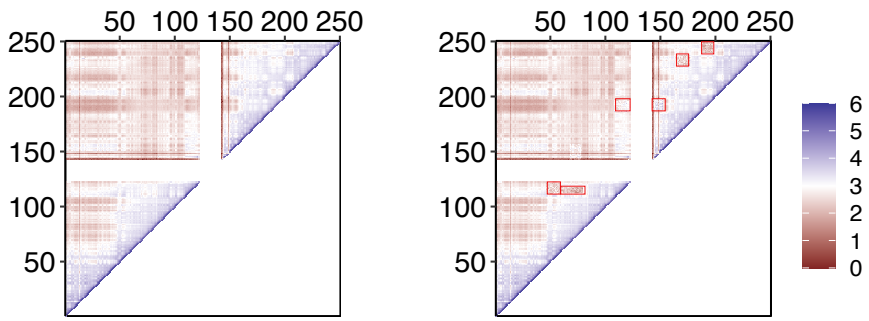

Figure S8: $\log_{10}$-transformed mean Hi-C interaction matrices (μ_X and μ_Y in Section S5.4) under the two conditions. DIR regions are highlighted in red squares on the right heatmap.

References

[1] Yong Zhang, Tao Liu, Clifford A Meyer, Jérôme Eeckhoute, David S Johnson, Bradley E Bernstein, Chad Nusbaum, Richard M Myers, Myles Brown, Wei Li, et al. Model-based analysis of chip-seq (macs). *Genome biology*, 9(9):1–9, 2008.

[2] Sven Heinz, Christopher Benner, Nathanael Spann, Eric Bertolino, Yin C Lin, Peter Laslo, Jason X Cheng, Cornelis Murre, Harinder Singh, and Christopher K Glass. Simple combinations of lineage-determining transcription factors prime cis-regulatory elements required for macrophage and b cell identities. *Molecular cell*, 38(4):576–589, 2010.

[3] David N Perkins, Darryl JC Pappin, David M Creasy, and John S Cottrell. Probability-based protein identification by searching sequence databases using mass spectrometry data. *ELECTROPHORESIS: An International Journal*, 20(18):3551–3567, 1999.

[4] Mark D Robinson, Davis J McCarthy, and Gordon K Smyth. edger: a bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, 26(1):139–140, 2010.

[5] Michael I Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for rna-seq data with deseq2. *Genome biology*, 15(12):550, 2014.

[6] Cole Trapnell, David G Hendrickson, Martin Sauvageau, Loyal Goff, John L Rinn, and Lior Pachter. Differential analysis of gene regulation at transcript resolution with rna-seq. *Nature biotechnology*, 31(1):46–53, 2013.

[7] Jun Li, Daniela M Witten, Iain M Johnstone, and Robert Tibshirani. Normalization, testing, and false discovery rate estimation for rna-sequencing data. *Biostatistics*, 13(3):523–538, 2012.

[8] Thomas J Hardcastle and Krystyna A Kelly. bayseq: empirical bayesian methods for identifying differential expression in sequence count data. *BMC bioinformatics*, 11(1):1–14, 2010.

[9] Gordon K Smyth. Linear models and empirical bayes methods for assessing differential expression in microarray experiments. *Statistical applications in genetics and molecular biology*, 3(1), 2004.

[10] Matthew E Ritchie, Belinda Phipson, DI Wu, Yifang Hu, Charity W Law, Wei Shi, and Gordon K Smyth. limma powers differential expression analyses for rna-sequencing and microarray studies. *Nucleic acids research*, 43(7):e47–e47, 2015.

[11] John C Stansfield, Kellen G Cresswell, and Mikhail G Dozmorov. multihiccompare: joint normalization and comparative analysis of complex hi-c experiments. *Bioinformatics*, 35(17):2916–2923, 2019.

[12] Mohamed Nadhir Djekidel, Yang Chen, and Michael Q Zhang. Find: differential chromatin interactions detection using a spatial poisson process. *Genome research*, 28(3):412–422, 2018.

[13] Aaron TL Lun and Gordon K Smyth. diffhic: a bioconductor package to detect differential genomic interactions in hi-c data. *BMC bioinformatics*, 16(1):1–11, 2015.

[14] 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.

[15] Bradley Efron and Robert Tibshirani. Empirical bayes methods and false discovery rates for microarrays. *Genetic epidemiology*, 23(1):70–86, 2002.

[16] Bradley Efron, Robert Tibshirani, John D Storey, and Virginia Tusher. Empirical bayes analysis of a microarray experiment. *Journal of the American statistical association*, 96(456):1151–1160, 2001.

[17] Matthew Stephens. False discovery rates: a new deal. *Biostatistics*, 18(2):275–294, 2017.

[18] John D Storey and Robert Tibshirani. Statistical significance for genomewide studies. *Proceedings of the National Academy of Sciences*, 100(16):9440–9445, 2003.

[19] Anat Reiner, Daniel Yekutieli, and Yoav Benjamini. Identifying differentially expressed genes using false discovery rate controlling procedures. *Bioinformatics*, 19(3):368–375, 2003.

[20] Bing Yang, Yan-Jie Wu, Ming Zhu, Sheng-Bo Fan, Jinzhong Lin, Kun Zhang, Shuang Li, Hao Chi, Yu-Xin Li, Hai-Feng Chen, et al. Identification of cross-linked peptides from complex samples. *Nature methods*, 9(9):904–906, 2012.

[21] James Robert White, Niranjan Nagarajan, and Mihai Pop. Statistical methods for detecting differentially abundant features in clinical metagenomic samples. *PLoS Comput Biol*, 5(4):e1000352, 2009.

[22] Andrey A Shabalin. Matrix eqtl: ultra fast eqtl analysis via large matrix operations. *Bioinformatics*, 28(10):1353–1358, 2012.

[23] Stijn Hawinkel, Federico Mattiello, Luc Bijnens, and Olivier Thas. A broken promise: microbiome differential abundance methods do not control the false discovery rate. *Briefings in bioinformatics*, 20(1):210–221, 2019.

[24] Ye Zheng and Sündüz Keleş. Freehi-c simulates high-fidelity hi-c data for benchmarking and data augmentation. *Nature Methods*, 17(1):37–40, 2020.

[25] John D Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.

[26] Yoav Benjamini and Yosef Hochberg. Multiple hypotheses testing with weights. *Scandinavian Journal of Statistics*, 24(3):407–418, 1997.

[27] Nikolaos Ignatiadis, Bernd Klaus, Judith B Zaugg, and Wolfgang Huber. Data-driven hypothesis weighting increases detection power in genome-scale multiple testing. *Nature methods*, 13(7): 577–580, 2016.

[28] Lihua Lei and William Fithian. Adapt: an interactive procedure for multiple testing with side information. *arXiv preprint arXiv:1609.06035*, 2016.

[29] Simina M Boca and Jeffrey T Leek. A direct approach to estimating false discovery rates conditional on covariates. *PeerJ*, 6:e6035, 2018.

[30] Wei Li, Han Xu, Tengfei Xiao, Le Cong, Michael I Love, Feng Zhang, Rafael A Irizarry, Jun S Liu, Myles Brown, and X Shirley Liu. Mageck enables robust identification of essential genes from genome-scale crispr/cas9 knockout screens. *Genome biology*, 15(12):554, 2014.

[31] Alexey I Nesvizhskii. A survey of computational methods and error rate estimation procedures for peptide and protein identification in shotgun proteomics. *Journal of proteomics*, 73(11):2092–2123, 2010.

[32] Lukas Käll, John D Storey, Michael J MacCoss, and William Stafford Noble. Assigning significance to peptides identified by tandem mass spectrometry using decoy databases. *Journal of proteome research*, 7(01):29–34, 2008.

[33] Christoffer Flensburg, Sarah A Kinkel, Andrew Keniry, Marnie E Blewitt, and Alicia Oshlack. A comparison of control samples for chip-seq of histone modifications. *Frontiers in genetics*, 5:329, 2014.

[34] Joses Ho, Tayfun Tumkaya, Sameer Aryal, Hyungwon Choi, and Adam Claridge-Chang. Moving beyond p values: data analysis with estimation graphics. *Nature methods*, 16(7):565–566, 2019.

[35] Abhishek K Sarkar and Matthew Stephens. Separating measurement and expression models clarifies confusion in single cell rna-seq analysis. *BioRxiv*, 2020.

[36] Rina Foygel Barber and Emmanuel J Candès. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085, 2015.

[37] Rina Foygel Barber, Emmanuel J Candès, et al. A knockoff filter for high-dimensional selective inference. *The Annals of Statistics*, 47(5):2504–2537, 2019.

[38] Jaime Roquero Gimenez and James Zou. Improving the stability of the knockoff procedure: Multiple simultaneous knockoffs and entropy maximization. *arXiv preprint arXiv:1810.11378*, 2018.

[39] Bradley Efron. Local false discovery rates, 2005.

[40] Xiaoquan Wen. A unified view of false discovery rate control: Reconciliation of bayesian and frequentist approaches. *arXiv preprint arXiv:1803.05284*, 2018.

[41] Peter J Park. Chip–seq: advantages and challenges of a maturing technology. *Nature reviews genetics*, 10(10):669–680, 2009.

[42] Pamela J Mitchell and Robert Tjian. Transcriptional regulation in mammalian cells by sequence-specific dna binding proteins. *Science*, 245(4916):371–378, 1989.

[43] Mark Ptashne and Alexander Gann. Transcriptional activation by recruitment. *Nature*, 386(6625): 569–577, 1997.

[44] Timothy Bailey, Pawel Krajewski, Istvan Ladunga, Celine Lefebvre, Qunhua Li, Tao Liu, Pedro Madrigal, Cenny Taslim, and Jie Zhang. Practical guidelines for the comprehensive analysis of chip-seq data. *PLoS Comput Biol*, 9(11):e1003326, 2013.

[45] ENCODE Project Consortium et al. An integrated encyclopedia of dna elements in the human genome. *Nature*, 489(7414):57–74, 2012.

[46] Aaron Diaz, Kiyoub Park, Daniel A Lim, and Jun S Song. Normalization, bias correction, and peak calling for chip-seq. *Statistical applications in genetics and molecular biology*, 11(3), 2012.

[47] Boris Bogdanow, Henrik Zauber, and Matthias Selbach. Systematic errors in peptide and protein identification and quantification by modified peptides. *Molecular & Cellular Proteomics*, 15(8): 2791–2801, 2016.

[48] Michael P Washburn, Dirk Wolters, and John R Yates. Large-scale analysis of the yeast proteome by multidimensional protein identification technology. *Nature biotechnology*, 19(3):242–247, 2001.

[49] Marshall Bern, Yong J Kil, and Christopher Becker. Byonic: advanced peptide and protein identification software. *Current protocols in bioinformatics*, 40(1):13–20, 2012.

[50] Claire R Williams, Alyssa Baccarella, Jay Z Parrish, and Charles C Kim. Empirical assessment of analysis workflows for differential expression analysis of human samples using rna-seq. *BMC bioinformatics*, 18(1):38, 2017.

[51] Keegan Korthauer, Patrick K Kimes, Claire Duvallet, Alejandro Reyes, Ayshwarya Subramanian, Mingxiang Teng, Chinmay Shukla, Eric J Alm, and Stephanie C Hicks. A practical guide to methods controlling false discoveries in computational biology. *Genome biology*, 20(1):1–21, 2019.

[52] Jonathan Thorsen, Asker Brejnrod, Martin Mortensen, Morten A Rasmussen, Jakob Stokholm, Waleed Abu Al-Soud, Søren Sørensen, Hans Bisgaard, and Johannes Waage. Large-scale benchmarking reveals false discoveries and count transformation sensitivity in 16s rrna gene amplicon data analysis methods used in microbiome studies. *Microbiome*, 4(1):62, 2016.

[53] Charlotte Soneson and Mark D Robinson. Bias, robustness and scalability in single-cell differential expression analysis. *Nature methods*, 15(4):255, 2018.

[54] Bradley Efron. Large-scale simultaneous hypothesis testing: the choice of a null hypothesis. *Journal of the American Statistical Association*, 99(465):96–104, 2004.

[55] Marina Spivak, Jason Weston, Léon Bottou, Lukas Käll, and William Stafford Noble. Improvements to the percolator algorithm for peptide identification from shotgun proteomics data sets. *Journal of proteome research*, 8(7):3737–3745, 2009.

[56] John C Stansfield, Kellen G Cresswell, Vladimir I Vladimirov, and Mikhail G Dozmorov. Hiccompare: an r-package for joint normalization and comparison of hi-c datasets. *BMC bioinformatics*, 19(1):279, 2018.

[57] Mark D Robinson and Alicia Oshlack. A scaling normalization method for differential expression analysis of rna-seq data. *Genome biology*, 11(3):1–9, 2010.