
Black Box FDR

Wesley Tansey^{1 2} Yixin Wang^{1 3} David M. Blei^{1 3 4} Raul Rabadan²

Abstract

Analyzing large-scale, multi-experiment studies requires scientists to test each experimental outcome for statistical significance and then assess the results as a whole. We present Black Box FDR (BB-FDR), an empirical-Bayes method for analyzing multi-experiment studies when many covariates are gathered per experiment. BB-FDR boosts power and controls the false discovery rate (FDR) at two stages of study analysis. In Stage 1, BB-FDR learns a deep model of the prior probability that an experiment will produce a significant outcome. It then uses posterior estimates to report discoveries at a given FDR threshold. In Stage 2, BB-FDR extracts insights from the overall study by determining which covariates have significant predictive power across all experiments. On a series of benchmarks, BB-FDR outperforms competing state-of-the-art methods in both stages of analysis. We apply BB-FDR to two real studies, each containing hundreds of experiments about cancer drug efficacy. For both studies, BB-FDR increases the proportion of significant outcomes discovered and selects variables that reveal key genomic drivers of drug sensitivity and resistance in cancer.

1. Introduction

High-throughput screening (HTS) techniques have fundamentally changed the landscape of modern biological experimentation. Rather than conducting just one experiment at a time, HTS enables scientists to perform hundreds of parallel experiments, each with different materials and different interventions. At the same time, HTS also enables scientists

¹Data Science Institute, Columbia University, New York, NY, USA ²Department of Systems Biology, Columbia University Medical Center, New York, NY, USA ³Department of Statistics, Columbia University, New York, NY, USA ⁴Department of Computer Science, Columbia University, New York, NY, USA. Correspondence to: Wesley Tansey <wt2274@columbia.edu>.

to gather rich contextual information about each experiment by profiling the materials under study using techniques like DNA sequencing. Thus, each HTS study produces a dataset of many experiments, where each experiment contains both an outcome variable and a high-dimensional feature set describing the context.

Figure 1 shows a slice of the Genomics of Drug Sensitivity in Cancer (GDSC) dataset (Yang et al., 2012), an HTS study investigating how cancer cell lines respond to different cancer therapeutics. The left panel shows the relative response of 30 different cancer cell lines (C_1, C_2, \dots, C_{30}) treated with the drug Nutlin-3. For each cell line, the treatment response (black triangles) is overlaid on top of the untreated control replicate distribution (gray box plots). Even when no drug is applied, each cell line still exhibits natural variation. The first goal in analyzing this data is therefore to address the question of whether a given cell line responded to the treatment. Concretely, we need to perform a hypothesis test for each cell line, where the null hypothesis is that the drug had no effect. Absent other information, this would be a classic *multiple hypothesis testing* (MHT) problem.

But HTS studies such as GDSC differ from the classical setup by also producing a rich set of side information for each experiment. The right panel of Figure 1 shows a subset of the genomic profile for each cell line, with a black dot indicating that the cell line has a mutation in that gene. Biologically, a mutated gene can lead to different phenotypic behavior that may cause different sensitivity or resistance to a drug.

Statistically, this means the likelihood of a cell line responding to treatment is a latent function of that cell line’s genomic profile. Identifying which mutations are associated with treatment response could guide future experiments and development of new targeted therapies. Deriving scientific insight from patterns across experiments represents a second phase of hypothesis testing, where the null hypothesis is that a given gene is not associated with drug response.

We term these two phases of hypothesis testing *Stage 1* and *Stage 2* analysis. This leads us to ask two scientifically-motivated statistical inference questions:

- **Stage 1:** How do we leverage the available side information in a HTS study to increase how many significant

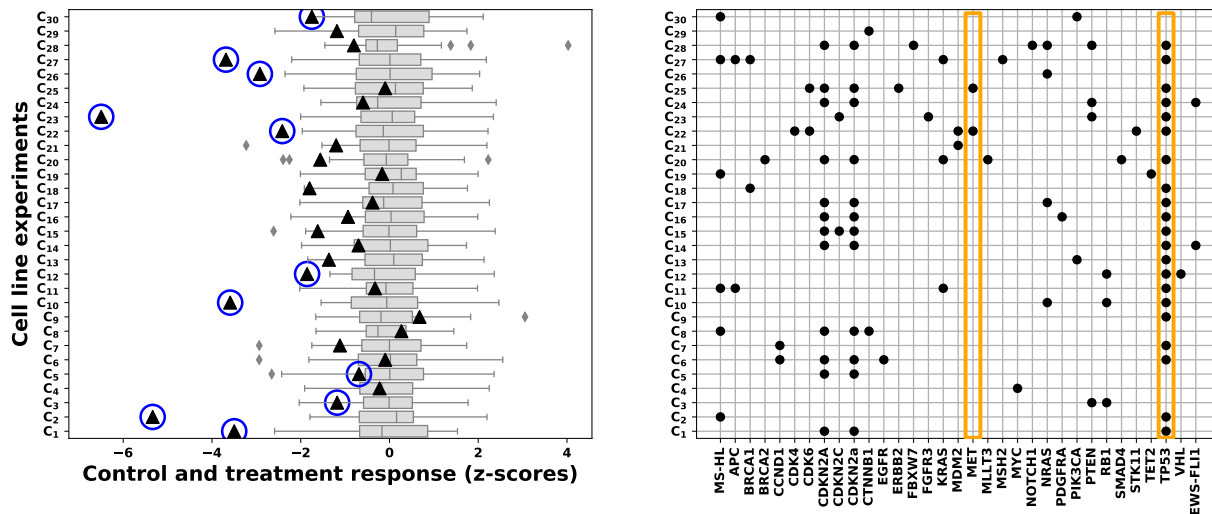


Figure 1. Left: a subset of 30 cell line experiments from the Nutlin-3 case study in Section 5. Control replicates (gray box plots) and cell line responses (black triangles) are measured as z -scores relative to mean control values. Right: a subset of the corresponding genomic features for each experiment; black dots indicate a cell line has a recurrent mutation in the gene labeled on the x-axis. The goal in Stage 1 analysis is to select cell lines that showed a significant response (circled in blue). In Stage 2, the genomic features are analyzed to understand the mutations driving drug response (circled in orange).

outcomes we can detect?

- **Stage 2:** Can we discover which variables are associated with significant outcomes, even when the underlying function is high-dimensional and nonlinear?

We answer both of these questions and propose Black Box FDR (BB-FDR), a method for analyzing multi-experiment studies with many covariates gathered per experiment. BB-FDR uses the covariates to build a deep probabilistic model that predicts how likely a given experiment is to generate significant outcomes *a priori*. It uses this prior model to adaptively select significant outcomes in a manner that controls the overall false discovery rate (FDR) at a specified Stage 1 level. BB-FDR then builds a predictive model of each covariate to perform variable selection on the Stage 1 model while conserving a specified Stage 2 FDR threshold.

We validate BB-FDR on both synthetic and real data. BB-FDR outperforms other state-of-the-art Stage 1 methods in a series of benchmarks, including the recently-proposed NeuralFDR (Xia et al., 2017). BB-FDR is also a more pragmatic choice compared to a fully-Bayesian approach: it scales trivially to thousands of covariates, can learn arbitrarily complex functions, and runs easily on a laptop. We apply BB-FDR to a real-world case study of two cancer drug screenings. BB-FDR finds more significant discoveries on the data and recovers an informative set of biologically-plausible genes that may convey drug sensitivity and resistance in cancer.

2. Multiple testing and FDR control

In the classical MHT setup, $\mathbf{z} = (z_1, \dots, z_n)$ are a set of independent observations of the outcomes of n experiments. For each experiment, a treatment is applied to a target and the treatment has either no effect ($h_i = 0$) or some effect ($h_i = 1$). If the treatment has no effect, the distribution of the test statistic is the null distribution $f_0(z)$; otherwise it follows an unknown alternative distribution $f_1(z)$. The null hypothesis for every experiment is that the test statistic was drawn from the null distribution: $H_0 : h_i = 0$.

2.1. False discovery rate control

In most experiments of interest, it is impossible to determine h_i with no error. For a given prediction \hat{h}_i , we say it is a true positive or a true discovery if $\hat{h}_i = 1 = h_i$ and a false positive or false discovery if $\hat{h}_i = 1 \neq h_i$. Let $\mathcal{S} = \{i : h_i = 1\}$ be the set of observations for which the treatment had an effect and $\hat{\mathcal{S}} = \{i : \hat{h}_i = 1\}$ be the set of predicted discoveries. We seek procedures that maximize the true positive rate (TPR) also known as *power*, while controlling the false discovery rate—the expected proportion of the predicted discoveries that are actually false positives,

$$\text{FDR} := \mathbb{E}[\text{FDP}], \quad \text{FDP} = \frac{\#\{i : i \in \hat{\mathcal{S}} \setminus \mathcal{S}\}}{\#\{i : i \in \hat{\mathcal{S}}\}}. \quad (1)$$

FDP in (1) is the *false discovery proportion*: the actual proportion of false positives in the predicted discovery set for a specific experiment. While ideally we would like to control the FDP, the randomness of the outcome variables

makes this impossible in practice. Thus FDR is the typical error measure targeted in modern scientific analyses.

2.2. Related work

Controlling FDR in multiple hypothesis testing has a long history in statistics and machine learning. The Benjamini-Hochberg (BH) procedure (Benjamini & Hochberg, 1995) is the classic technique and still the most widely used in science. Many other methods have since been developed to take advantage of study-specific information to increase power. Recent examples include accumulation tests for ordering information (Li & Barber, 2017), the p-filter for grouping and test statistic dependency (Ramdas et al., 2017), FDR smoothing for spatial testing (Tansey et al., 2017), FDR-regression for low-dimensional covariates (Scott et al., 2015), and, most recently, NeuralFDR for high-dimensional covariates (Xia et al., 2017). We consider high-dimensional covariates and compare against NeuralFDR in Section 4.

3. Black Box FDR

Consider a study with n independent experiments that produces a set of independent test statistics $\mathbf{z} = (z_1, \dots, z_n)$ corresponding to the outcome measurements, as in Section 2. However, now each experiment also has a vector of m covariates $X_i = (X_{i1}, \dots, X_{im})$ containing side information that may influence the outcome of that experiment. Specifically, whether the experiment comes from the null distribution $h_i = 0$ or the alternative $h_i = 1$ is allowed to depend arbitrarily on X_i .

BB-FDR extends the empirical-Bayes two-groups model of Efron (2008) by building a hierarchical probabilistic model with experiment-specific priors modeled through a deep neural network. We first estimate the alternative distribution offline using predictive recursion (Newton, 2002) to estimate f_1 . This follows other recent extensions to the two-groups model (Scott et al., 2015; Tansey et al., 2017) and enjoys strong empirical performance and consistency guarantees (Tokdar et al., 2009). BB-FDR then focuses on modeling the experiment-specific prior, assuming the null and alternative distributions are fixed.

3.1. Stage 1: determining significant outcomes

We model the test statistic as arising from a mixture model of two components, the null (f_0) and the alternative (f_1). An experiment-specific weight c_i then models the prior probability of the test statistic coming from the alternative (i.e. the probability of the treatment having an effect *a priori*). We place a beta prior on each experiment-specific prior c_i and model the parameters of the hyperprior with a black box function G parameterized by θ ; in our implementation, G is




Figure 2. The graphical model for BB-FDR.

a deep neural network. The complete model for BB-FDR is:

$$\begin{aligned} z_i &\sim h_i f_1(z_i) + (1 - h_i) f_0(z_i) \\ h_i &\sim \text{Bernoulli}(c_i) \\ c_i &\sim \text{Beta}(a_i, b_i) \\ a_i, b_i &= G_{\theta, i}(X) \end{aligned} \quad (2)$$

We optimize θ by integrating out h_i and maximizing the complete data log-likelihood,

$$p_{\theta}(z_i) = \int_0^1 (c_i f_1(z_i) + (1 - c_i) f_0(z_i)) \text{Beta}(c_i | G_{\theta, i}(X)) dc_i. \quad (3)$$

Figure 2 shows the BB-FDR graphical model.

The beta prior is a departure from other two-groups extensions, which use a flatter hierarchy and learn a predictive model for c_i (Scott et al., 2015; Tansey et al., 2017). We found the flat approach to be difficult to train, leading to a degenerate G that always predicts the global mean prior.

A hierarchical prior improves training for two reasons. First, optimization is easier and more stable because the output of the function is two soft-plus activations. Compared to a sigmoid, this form leads to less saturated gradients. Second, the additional hierarchy allows the model to assign different degrees of confidence to each experiment, changing the model from homoskedastic to heteroskedastic. Finally, we found it important to enforce concavity of the beta distribution; we thus add 1 to both a_i and b_i .

We fit the model in (2) with stochastic gradient descent (SGD) on an L_2 -regularized loss function,

$$\underset{\theta \in \mathcal{R}^{|\theta|}}{\text{minimize}} \quad - \sum_i \log p_{\theta}(z_i) + \lambda \|G_{\theta}(X)\|_F^2, \quad (4)$$

where $\|\cdot\|_F$ is the Frobenius norm. In pilot studies, we found adding a small amount of L_2 -regularization prevented over-fitting at virtually no cost to statistical power. For computational purposes, we approximate the integral in (3) by a fine-grained numerical grid.

Once the optimized parameters $\hat{\theta}$ are chosen, we calculate the posterior probability of each test statistic coming from

the alternative,

$$\begin{aligned} \hat{w}_i &= p_{\hat{\theta}}(h_i = 1 | z_i) \\ &= \int_0^1 \frac{c_i f_1(z_i) \text{Beta}(c_i | G_{\hat{\theta}, i}(X))}{c_i f_1(z_i) + (1 - c_i) f_0(z_i)} dc_i. \end{aligned} \quad (5)$$

Assuming the posteriors are accurate, rejecting the i^{th} hypothesis will produce $1 - \hat{w}_i$ false positives in expectation. Therefore we can maximize the total number of discoveries by a simple step down procedure. First, we sort the posteriors in descending order by the likelihood of the test statistics being drawn from the alternative. We then reject the first q hypotheses, where $0 \leq q \leq n$ is the largest possible index such that the expected proportion of false discoveries is below the FDR threshold. Formally, this procedure solves the optimization problem,

$$\begin{aligned} &\underset{q}{\text{maximize}} && q \\ &\text{subject to} && \frac{\sum_{i=1}^q (1 - \hat{w}_i)}{q} \leq \alpha, \end{aligned} \quad (6)$$

for a given FDR threshold α . By convention $\frac{0}{0} = 0$.

The model in (2)–(6) handles Stage 1 of the analysis. The black box model G uses the entire feature vector X_i of every experiment to predict the prior parameters over c_i . The observations z_i are then used to calculate the posterior probability \hat{w}_i that the treatment had an effect. The selection procedure in (6) uses these posteriors to reject a maximum number of null hypotheses while conserving the FDR.

3.2. Stage 2: identifying important variables

Using a flexible black box model for G in (2) provides a trade-off. On one hand, it enables BB-FDR to learn a rich class of functions for the relationship between the covariates and the test statistic. As we show in Section 4, this increases power in Stage 1 compared to a standard linear model.

However, variable selection (Stage 2) is straightforward in linear models whereas black box models are by definition opaque. Understanding which variables deep learning models use to make predictions is an ongoing area of research in both machine learning (e.g. Elenberg et al., 2017) and specific scientific disciplines (e.g. Olden & Jackson, 2002; Giam & Olden, 2015, in ecology). As far as we are aware, there are currently no methods that provide variable selection with FDR control when the covariates may have arbitrary dependency structure.

To overcome this challenge, BB-FDR uses conditional randomization tests (CRTs) (Candes et al., 2018). The idea of a CRT is to model each coordinate of the feature matrix $X_{\cdot j}$ using only the other coordinates $X_{\cdot -j}$. The conditional distribution $\mathbb{P}(X_{\cdot j} | X_{\cdot -j})$ then represents a valid null distribution for testing the hypothesis $X_{\cdot j} \perp\!\!\!\perp Z | X_{\cdot -j}$, where

Z is the test statistic. The corresponding p -value can be calculated by sampling from the conditional to approximate the true p -value,

$$p_j = \mathbb{E}_{\tilde{X}_{\cdot j} \sim \mathbb{P}(X_{\cdot j} | X_{\cdot -j})} \left[\mathbb{I} \left[t(\mathbf{z}, X) \leq t(\mathbf{z}, (\tilde{X}_{\cdot j}, X_{\cdot -j})) \right] \right],$$

where t is the test statistic of interest. Once the p -values have been estimated for all features, we can apply standard BH correction and report significant features.

BB-FDR tests which features are associated with a change in the posterior probability of z_i coming from the alternative. It uses the negative entropy of the posteriors as the test statistic,

$$t(\mathbf{z}, X) = \sum_i \hat{w}_i \log \hat{w}_i + \sum_i (1 - \hat{w}_i) \log(1 - \hat{w}_i). \quad (7)$$

Intuitively, if a feature is useful in predicting treatment efficacy, it should reduce the overall entropy of the posterior. By definition, a feature sampled from the null adds no new information to the model; it cannot systematically reduce the entropy.

For this procedure to retain frequentist consistency guarantees, both the conditional null distribution and the model of the prior must be the true distributions. In practice, one never has access to these and thus we estimate both; for the conditional null, we use gradient boosting trees (Chen & Guestrin, 2016).

4. Benchmarks

We perform a series of benchmark studies to assess the performance of BB-FDR in both stages of inference. For each benchmark, we compare the power of BB-FDR to other state-of-the-art approaches. In all studies, we consider binary covariates and real-valued z-scores as test statistics.

Across experiments, we found BB-FDR is particularly suitable for large samples: it outperforms competing methods in both stages while being more computationally efficient.

4.1. Setup

We consider three different ground truth models for $P(X)$, the joint distribution over the covariates, and $P(h = 1 | X)$, the prior probability of coming from the alternative distribution given the covariates:

- **Constant:** All covariates are sampled IID normal; the prior is independent of the covariates, with $P(h_i = 1 | X) = 0.5$.
- **Linear:** Covariates are sampled from a multivariate normal with full covariance matrix (i.e. conditionally linear); the prior is a linear function with IID standard normal coefficients for each covariate.

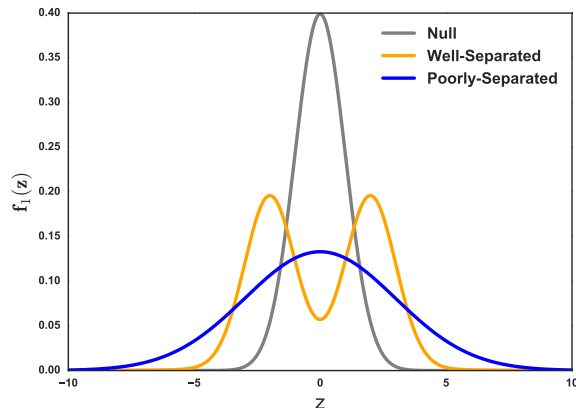


Figure 3. The two alternative densities used in our benchmarks. The well-separated (WS) density has little overlap with the null, making for a stronger signal.

- **Nonlinear:** Covariates and prior coefficients are generated similarly. We first drawing 20 IID uniform Bern(0.5) latent variables. For each covariate, 5 pairs of latent variables (u_i, u_j) are chosen and with equal probability are either ANDed or XORed together and multiplied by a draw from a standard normal; the latent weights are summed to get the final logit value for the covariate or coefficient.

For each of the three ground truth models, we consider two different alternative distributions:

- **Well-Separated (WS):** A 3-component Gaussian mixture model, $f_1(z) = 0.48\mathcal{N}(-2, 1) + 0.04\mathcal{N}(0, 16) + 0.48\mathcal{N}(2, 1)$
- **Poorly-Separated (PS):** A single normal with high overlap with the null, $f_1(z) = \mathcal{N}(0, 9)$.

Figure 3 shows the densities used in our benchmarks.

For each of the 6 combinations of the above scenarios, we run 100 independent trials. Each trial uses 50 covariates; for all trials with a non-constant prior, 25 of the variables are used in the true data generating distribution and the other 25 are null variables with no association with the outcome. To measure sample efficiency, we also vary the sample size from $n = 100$ to $n = 50K$. The target FDR threshold is set to 10% for both stages of inference.

We compare BB-FDR to the classic Benjamini-Hochberg (BH) method (Benjamini & Hochberg, 1995), the recently-proposed NeuralFDR (Xia et al., 2017), and a fully-Bayesian logistic regression model for c_i in place of the black box prior in (2). For NeuralFDR, we use the default recommended settings, including five random restarts and a ten-layer deep neural network. The fully-Bayesian

method uses a standard normal prior on the weights and an inverse-Wishart prior on the variance, with weak hyperpriors. In the nonlinear scenario, we specify all possible pairwise interactions as the covariate set for the fully-Bayesian model to ensure it is well-specified. We fit the model using Polya-gamma sampling (Polson et al., 2013) with 5000 burn-in iterations and 1000 samples. For BB-FDR, we use a $50 \times 200 \times 200 \times 2$ network with ReLU activation; for training we use RMS-prop (Tieleman & Hinton, 2012) with dropout, learning rate 3×10^{-4} , and batch size 100, and run for 50 epochs, with 3 folds to create 3 separate models as in NeuralFDR; we set the λ regularization term to 10^{-4} .

4.2. Stage 1 performance

Figure 4 shows the results for the Stage 1 benchmarks, where the goal is to determine for which experiments the treatment had an effect. The four methods generally conserve FDR at the specified 10% threshold, though NeuralFDR seems to systematically violate FDR in the low-sample regime.

Across all experiments, we see that both BH and NeuralFDR under-perform the two Bayesian methods. In the case of BH, this is straight-forward as it uses only the p-value from each experiment and has no notion of side information. NeuralFDR, on the other hand, uses a deep neural network and several random restarts. There are a few possible reasons for its poor performance. First, the NeuralFDR method was reported to be very difficult to train by the original authors, so it is possible that it is simply not finding good fits of the model. Second, BB-FDR assumes that the alternative distribution is conditionally independent of the prior; NeuralFDR makes no such assumption and may lose some power as a result. Finally, NeuralFDR was originally tested on 1- and 2-dimensional problems against relatively weak baselines. Our benchmarks examine its performance in a higher-dimensional setting and with several uninformative features that may make fitting NeuralFDR difficult.

Since the fully-Bayesian method is well-specified in every benchmark, it serves as an oracle model to establish a reasonable upper bound on Stage 1 performance. However, the oracle power depends on the MCMC approximation of the posterior being well-mixed. As the sample size grows, the empirical-Bayes model used by BB-FDR gains an increasingly precise approximation to the true posterior. In the large-sample regime with a well-separated alternative, BB-FDR outperforms even the oracle. Furthermore, the fully-Bayesian method takes several hours to fit in the nonlinear scenarios; BB-FDR fits within a few minutes and can easily be run on a laptop.

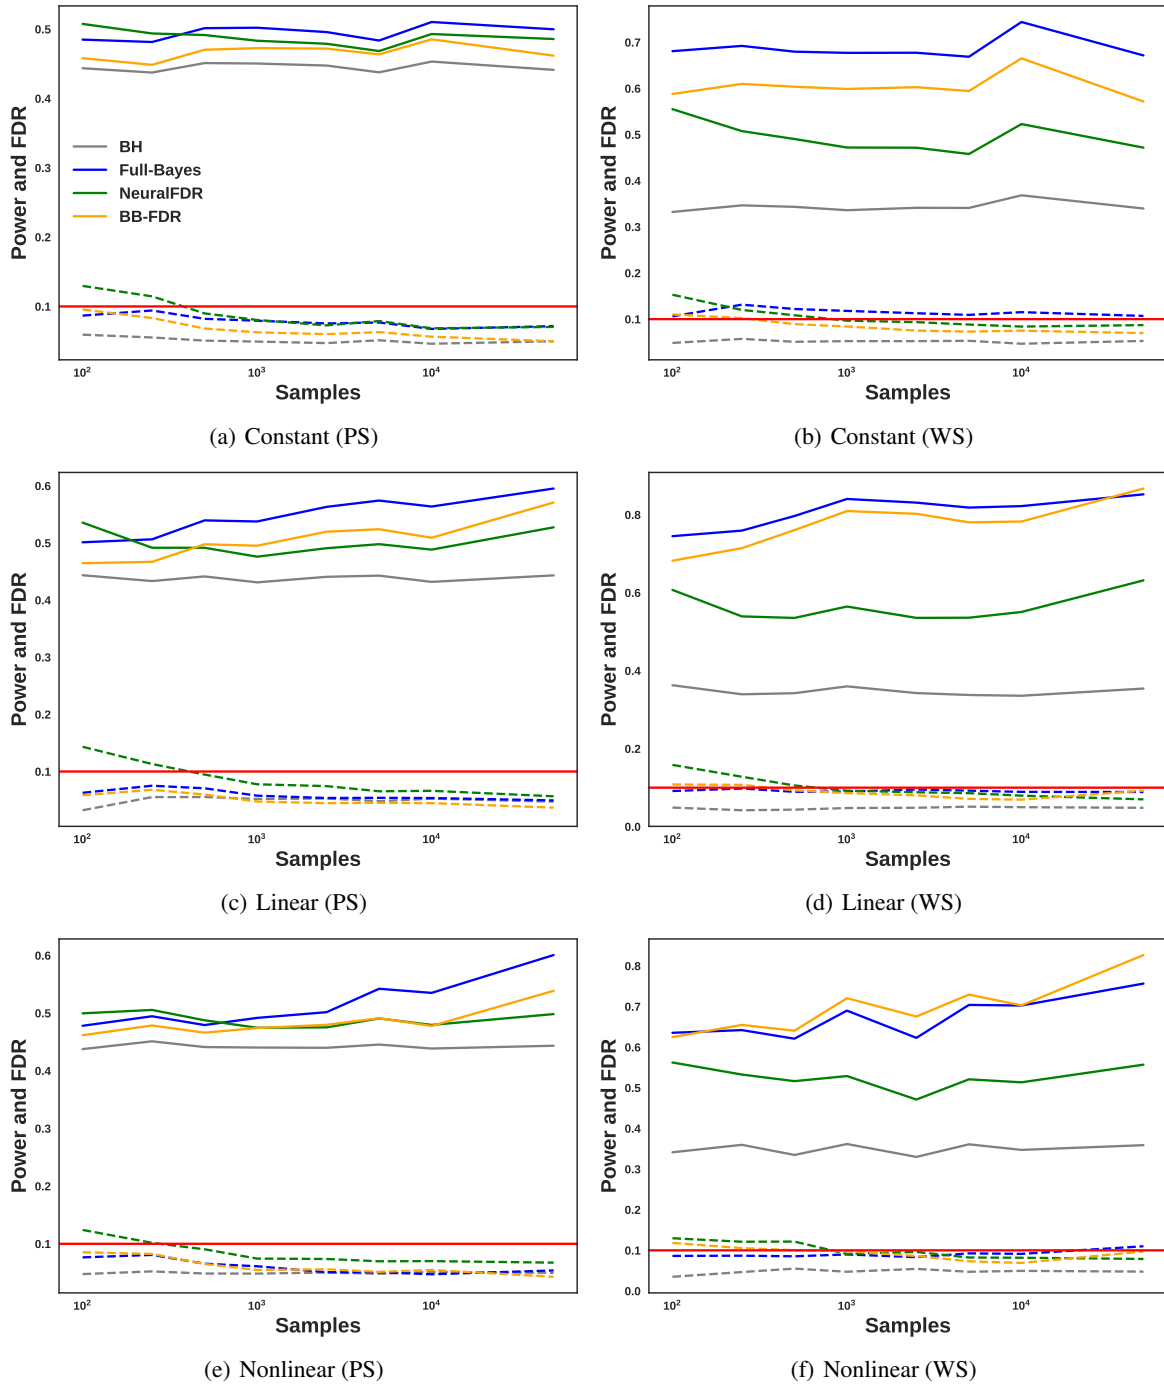


Figure 4. Hypothesis testing results on the synthetic datasets averaged over 100 trials at varying sample sizes on the two different alternative distributions. Solid lines show power; dashed lines show estimated FDR; the red horizontal line denotes the specified 10% FDR threshold. In general, the Benjamini-Hochberg and NeuralFDR methods have lower power since they do not model the alternative. The fully-Bayesian method has high power in the low-to-moderate sample regime, but as the sample size grows the empirical-Bayes approach of BB-FDR becomes more effective.




Figure 5. Variable selection results at a 10% FDR threshold. In low sample regimes, the conditional null distribution used in the CRT procedure is poorly fit and results in violated FDR thresholds. At moderate-to-large samples, BB-FDR has higher power than the fully-Bayesian model and conserves FDR.

4.3. Stage 2 performance

Neither BH nor NeuralFDR provide support for detecting important features (Stage 2), so we could not compare against them. For the Bayesian linear regression, we take the 90% posterior credible interval over the coefficient value for each covariate. If the interval does not contain zero, we reject the null hypothesis and report it as a discovered feature; this approach is standard in the Bayesian literature (Gelman et al., 2014).

Figure 5 presents the results of the variable selection benchmarks for the poorly-separated alternative distribution; results for the well-separated are similar. We omit the constant scenario, since there are no features to discover. In the small sample regime, the conditional distributions are poor estimators of the conditional null distribution for each feature. This leads to BB-FDR overestimating the number of signal features and violating the FDR threshold. As the sample size grows, the conditional null and black box prior become more accurate, leading to FDR control and higher power, respectively; in the large-sample regime, BB-FDR outperforms the fully-Bayesian approach.

We conclude by noting that BB-FDR is competitive with the fully-Bayesian approach even when the latter is well-specified. In practical data analysis scenarios, such as the cancer study we discuss next, we do not know the true prior function. It may easily contain many higher-order interaction terms that are prohibitive to consider explicitly in a fully-Bayesian model, making BB-FDR a pragmatic choice for real-world scientific datasets.

5. Cancer drug screening

As a case study of how BB-FDR is useful in practice, we apply it to two high-throughput cancer drug screening studies (Lapatinib and Nutlin-3) from the Genomics of Drug Sensitivity in Cancer (GDSC) (Yang et al., 2012). For both drug studies, BB-FDR increases the number of Stage 1 discoveries over classical BH correction; results on NeuralFDR were similar to BH and are omitted. In Stage 2, BB-FDR discovers biologically-plausible genes that may have a causal link to drug sensitivity and resistance. Experimental details are available in the supplement.

5.1. Analysis overview

Analysis of the two drug studies broadly follows the two stages outlined in the motivating example in Section 1. The Stage 1 task is to determine, for a specific drug being tested on a specific cell line, whether the drug had any effect. As with any biological process, natural variation injects randomness at many levels of the experiment: how fast the cells grow, how each cell responds, etc. Thus Stage 1 requires performing statistical hypothesis testing to determine if the cell population after treatment is truly smaller than would be expected from a control (null) population.

The inferential goal in Stage 2 is to gain scientific insight about which genes may be driving drug response. This involves building a statistical model of the relationship between the genomic profile of a cell line and its corresponding treatment response, then performing variable selection on the model. The selected genes form the basis for potential mechanisms of action and future experiments can be designed to test for a causal link or to investigate new drugs that better-target the proteins encoded by the genes.

| Lapatinib | Nutlin-3 |
|----------------------------------------|------------------------------------------------------------|
| BRCA1, BRCA2, CDK4 FGFR2, KIT, MSH2 | P300, FLCN, FLT3 MET, KIT, MSH6 SETD2, TP53, BCR-ABL |

Table 1. Significant gene mutations identified by BB-FDR that are associated with sensitivity and resistance to each drug. Both lists align well with known genomic targets of Lapatinib and Nutlin-3.

5.2. Results

Figure 6 shows the aggregate number of treatment effects discovered by both BH and BB-FDR. For both drugs, BB-FDR provides approximately a 50% increase in Stage 1 discoveries compared to BH. The genomic profiles of the cell lines provide enough prior information that even some outcomes with a z -score above zero are still found to be significant. This flexibility is impossible with classical Stage 1 testing methods like BH that do not consider covariate information.

Table 1 lists the genes reported by BB-FDR in Stage 2. Interpreting the quality of the results requires familiarity with genomics and cancer biology. Below, we briefly detail the scientific rationale behind the biological plausibility of the Stage 2 results and refer the reader to Weinberg (2013) for a full review.

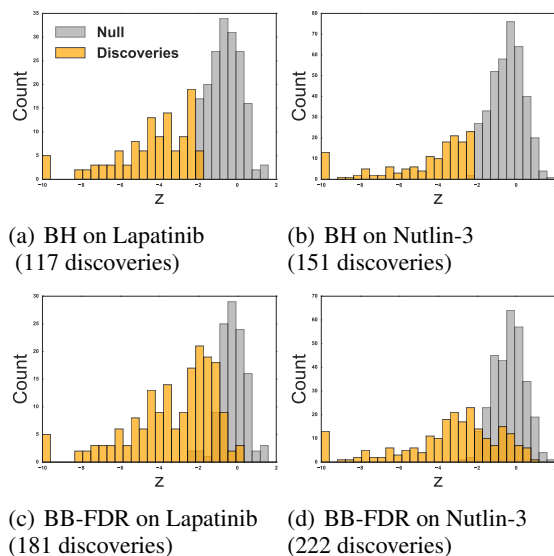


Figure 6. Discoveries found by BB-FDR on the two drugs, compared to the discoveries found by a naive BH approach. BB-FDR leverages the genomic profiling information of the cell lines to identify $\approx 50\%$ more discoveries at the same 10% FDR threshold.

Lapatinib has been approved for the treatment of HER2-positive breast cancers. BB-FDR indicates that BRCA1 and BRCA2 are associated with responses to Lapatinib. Both are

tumor suppressor genes that are seen mutated in more than 10% of breast cancers (BRCA stands for “breast cancer”) and thus cancer type may represent a latent confounder for drug efficacy that induces a conditional dependence. Lapatinib targets over-expression of the gene ERBB2 which can be caused by a mutant CDK4 gene. FGFR2 and KIT are also commonly associated with breast cancers (Slattery et al., 2013; Zhu et al., 2014) and BRCA1 is known to induce inactivation of the tumor suppressor MSH2 (Atalay et al., 2002). Given Lapatinib’s success as a breast cancer drug, the connection between all of the selected genes and breast cancer is a reassuring sign that BB-FDR selected biologically plausible features.

Nutlin-3 is an inhibitor of the oncogene MDM2, which negatively-regulates TP53. When highly over-expressed, MDM2 can functionally inactivate TP53. By targeting MDM2, Nutlin-3 enables a non-mutated (“wild type”) TP53 to trigger apoptosis in cancer cells. However, if TP53 is mutated, Nutlin-3 will be ineffective and hence its mutation state is an important driver of Nutlin-3 sensitivity. When wild type TP53 is present, MET controls the fate of the cell (Sullivan et al., 2012), SETD2 functionality is required to activate TP53 (Carvalho et al., 2014), P300 mediates TP53 acetylation (Reed & Quelle, 2014), and BCR-ABL is a gene fusion that induces loss of TP53 (Pierce et al., 2000). These genes interact in complex, non-linear ways, yet BB-FDR is still able to identify them as important. Finally, FLCN is a tumor suppressor gene that can delay cell cycle like TP53 (Laviolette et al., 2013). The mechanism by which FLCN and TP53 are interrelated is currently unclear, representing a potential target for future experiments.

6. Conclusion

We presented Black Box FDR (BB-FDR), an empirical-Bayes method that increases statistical power in multi-experiment scientific studies when side information is available for each experiment. BB-FDR combines deep probabilistic modeling with recent multiple testing techniques to boost testing power without sacrificing interpretability. Benchmarks show that BB-FDR outperforms state-of-the-art techniques, often substantially and under a wide array of experimental conditions. BB-FDR also finds more experimental discoveries two cancer drug screening datasets and provides scientific insight into the mechanisms associated with differential treatment response in cancer.

This work was supported by a pilot grant from Columbia University, NIH U54 CA193313, ONR N00014-15-1-2209, ONR 133691-5102004, NIH 5100481-5500001084, NSF CCF-1740833, the Alfred P. Sloan Foundation, the John Simon Guggenheim Foundation, Facebook, Amazon, and IBM.

References

- Atalay, Arzu, Crook, Tim, Ozturk, Mehmet, and Yulug, Isik G. Identification of genes induced by *brca1* in breast cancer cells. *Biochemical and biophysical research communications*, 299(5):839–846, 2002.
- Benjamini, Yoav and Hochberg, Yosef. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pp. 289–300, 1995.
- Candes, Emmanuel, Fan, Yingying, Janson, Lucas, and Lv, Jinchu. Panning for gold: model-xknockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2018.
- Carvalho, Sílvia, Vítor, Alexandra C, Sridhara, Sreerama C, Martins, Filipa B, Raposo, Ana C, Desterro, Joana MP, Ferreira, João, and de Almeida, Sérgio F. Setd2 is required for dna double-strand break repair and activation of the p53-mediated checkpoint. *Elife*, 3, 2014.
- Chen, Tianqi and Guestrin, Carlos. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794. ACM, 2016.
- Efron, Bradley. Microarrays, empirical bayes and the two-groups model. *Statistical science*, pp. 1–22, 2008.
- Elenberg, Ethan, Dimakis, Alexandros G, Feldman, Moran, and Karbasi, Amin. Streaming weak submodularity: Interpreting neural networks on the fly. In *Advances in Neural Information Processing Systems*, pp. 4047–4057, 2017.
- Gelman, Andrew, Carlin, John B, Stern, Hal S, Dunson, David B, Vehtari, Aki, and Rubin, Donald B. *Bayesian data analysis*, volume 2. CRC press Boca Raton, FL, 2014.
- Giam, Xingli and Olden, Julian D. A new r^2 -based metric to shed greater insight on variable importance in artificial neural networks. *Ecological Modelling*, 313:307–313, 2015.
- Laviolette, Laura A, Wilson, Jonas, Koller, Julia, Neil, Christopher, Hulick, Peter, Rejtar, Tomas, Karger, Barry, Teh, Bin Tean, and Iliopoulos, Othon. Human folliculin delays cell cycle progression through late s and g2/m-phases: effect of phosphorylation and tumor associated mutations. *PLoS one*, 8(7):e66775, 2013.
- Li, Ang and Barber, Rina Foygel. Accumulation tests for fdr control in ordered hypothesis testing. *Journal of the American Statistical Association*, 112(518):837–849, 2017.
- Newton, Michael A. A nonparametric recursive estimator of the mixing distribution. *Sankhya, Series A*, 64:306–22, 2002.
- Olden, Julian D and Jackson, Donald A. Illuminating the black box: a randomization approach for understanding variable contributions in artificial neural networks. *Ecological modelling*, 154(1-2):135–150, 2002.
- Pierce, Andrew, Spooncer, Elaine, Wooley, Sarah, Dive, Caroline, Francis, Julia M, Miyan, Jaleel, Owen-Lynch, P Jane, Dexter, T Michael, and Whetton, Anthony D. Bcr-abl protein tyrosine kinase activity induces a loss of p53 protein that mediates a delay in myeloid differentiation. *Oncogene*, 19(48):5487, 2000.
- Polson, Nicholas G, Scott, James G., and Windle, Jesse. Bayesian inference for logistic models using p γ -gamma latent variables. *Journal of the American statistical Association*, 108(504):1339–1349, 2013.
- Ramdas, Aaditya, Barber, Rina Foygel, Wainwright, Martin J, and Jordan, Michael I. A unified treatment of multiple testing with prior knowledge. *arXiv preprint arXiv:1703.06222*, 2017.
- Reed, Sara M and Quelle, Dawn E. p53 acetylation: regulation and consequences. *Cancers*, 7(1):30–69, 2014.
- Scott, James G, Kelly, Ryan C, Smith, Matthew A, Zhou, Pengcheng, and Kass, Robert E. False discovery rate regression: an application to neural synchrony detection in primary visual cortex. *Journal of the American Statistical Association*, 110(510):459–471, 2015.
- Slattery, Martha L, John, Esther M, Stern, Mariana C, Herrick, Jennifer, Lundgreen, Abbie, Giuliano, Anna R, Hines, Lisa, Baumgartner, Kathy B, Torres-Mejia, Gabriela, and Wolff, Roger K. Associations with growth factor genes (*fgf1*, *fgf2*, *pdgfb*, *fgfr2*, *nrg2*, *egf*, *erbb2*) with breast cancer risk and survival: the breast cancer health disparities study. *Breast cancer research and treatment*, 140(3):587–601, 2013.
- Sullivan, Kelly D, Padilla-Just, Nuria, Henry, Ryan E, Porter, Christopher C, Kim, Jihye, Tentler, John J, Eckhardt, S Gail, Tan, Aik Choon, DeGregori, James, and Espinosa, Joaquín M. Atm and met kinases are synthetic lethal with nongenotoxic activation of p53. *Nature chemical biology*, 8(7):646, 2012.
- Tansey, Wesley, Koyejo, Oluwasanmi, Poldrack, Russell A., and Scott, James G. False discovery rate smoothing. *Journal of the American Statistical Association*, (just-accepted), 2017.

- Tieleman, Tijmen and Hinton, Geoffrey. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2):26–31, 2012.
- Tokdar, Surya, Martin, R., and Ghosh, J.K. Consistency of a recursive estimate of mixing distributions. *The Annals of Statistics*, 37(5A):2502–22, 2009.
- Weinberg, Robert. *The biology of cancer*. Garland science, 2013.
- Xia, Fei, Zhang, Martin J, Zou, James Y, and Tse, David. Neuralfdr: Learning discovery thresholds from hypothesis features. In *Advances in Neural Information Processing Systems*, pp. 1540–1549, 2017.
- Yang, Wanjuan, Soares, Jorge, Greninger, Patricia, Edelman, Elena J, Lightfoot, Howard, Forbes, Simon, Bindal, Nidhi, Beare, Dave, Smith, James A, Thompson, I Richard, et al. Genomics of drug sensitivity in cancer (gdsc): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic acids research*, 41(D1):D955–D961, 2012.
- Zhu, Yun, Wang, Yan, Guan, Bing, Rao, Qiu, Wang, Jiandong, Ma, Henghui, Zhang, Zhihong, and Zhou, Xiaojun. C-kit and pdgfra gene mutations in triple negative breast cancer. *International journal of clinical and experimental pathology*, 7(7):4280, 2014.