Removing batch effects for prediction problems with frozen surrogate variable analysis

Hilary S. Parker[1], Héctor Corrada Bravo[2] and Jeffrey T. Leek[1]

[1] Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, Baltimore, MD, USA
[2] Center for Bioinformatics and Computational Biology, Department of Computer Science, University of Maryland, College Park, MD, USA

ABSTRACT

Batch effects are responsible for the failure of promising genomic prognostic signatures, major ambiguities in published genomic results, and retractions of widely-publicized findings. Batch effect corrections have been developed to remove these artifacts, but they are designed to be used in population studies. But genomic technologies are beginning to be used in clinical applications where samples are analyzed one at a time for diagnostic, prognostic, and predictive applications. There are currently no batch correction methods that have been developed specifically for prediction. In this paper, we propose an new method called frozen surrogate variable analysis (fSVA) that borrows strength from a training set for individual sample batch correction. We show that fSVA improves prediction accuracy in simulations and in public genomic studies. fSVA is available as part of the sva Bioconductor package.

Subjects Bioinformatics, Computational Biology, Statistics, Computational Science
Keywords Batch effects, Surrogate variable analysis, Prediction, Machine learning, Database, Statistics, Genomics

INTRODUCTION

Genomic technologies were originally developed and applied for basic science research and hypothesis generation (*Eisen et al., 1998*). As these technologies mature, they are increasingly being used as clinical tools for diagnosis or prognosis (*Chan & Ginsburg, 2011*). The high-dimensional measurements made by microarrays can be used to classify patients into predictive, prognostic, or diagnostic groups. Despite the incredible clinical promise of these technologies there have only been a few signatures that have successfully been translated into the clinic.

One of the reasons for the relatively low rate of success is the impact of unmeasured technological or biological confounders. These artifacts are collectively referred to as "batch effects" because the processing date, or batch, is the most commonly measured surrogate for these unmeasured variables in genomic studies (*Scharpf et al., 2011*; *Johnson, Li & Rabinovic, 2007*; *Walker et al., 2008*). The umbrella term batch effects also refers to any unmeasured variables that can vary from experiment to experiment, ranging from the technician who performs the experiment to the temperature and ozone levels that day (*Lander, 1999*; *Fare et al., 2003*).

Batch effects are responsible for the failure of promising genomic prognostic signatures (*Baggerly, Morris & Coombes, 2004*; *Baggerly, Coombes & Morris, 2005*), major ambiguities in published genomic results (*Spielman et al., 2007*; *Akey et al., 2007*), and retractions of widely-publicized findings (*Sebastiani et al., 2011*; *Lambert & Black, 2012*). In many experiments, the signal from these unmeasured confounders is larger than the biological signal of interest (*Leek et al., 2010*). But the impact of batch effects on prediction problems has only recently been demonstrated (*Parker & Leek, 2012*; *Luo et al., 2010*). Batch effects were also recognized as a significant hurdle in the development of personalized genomic biomarkers in the Institute of Medicine's report on clinical genomics (*Micheel, Nass & Omenn, 2012*).

While a number of methods have been developed for removing batch effects in population-based genomic studies (*Johnson, Li & Rabinovic, 2007*; *Gagnon-Bartsch & Speed, 2012*; *Leek & Storey, 2007*; *Leek et al., 2010*; *Walker et al., 2008*), there is currently no method for removing batch effects for prediction problems. There are two key differences between population level corrections and corrections designed for prediction problems. First, population level corrections assume that the biological groups of interest are known in advance. In prediction problems, the goal is to predict the biological group. Second, in prediction problems, new samples are observed one at a time, so the surrogate batch variable will have a unique and unknown value for each sample.

Here we propose frozen Surrogate Variable Analysis (fSVA) as a method for batch correction in prediction problems. fSVA borrows strength from a reference database to address the challenges unique to batch correction for prediction. The fSVA approach has two main components. First, surrogate variable analysis (SVA) is used to correct for batch effects in the training database. Any standard classification algorithm can then be applied to build a classifier based on this clean training data set. Second, probability weights and coefficients estimated on the training database are used to remove batch effects in new samples. The classifier trained on the clean database can then be applied to these cleaned samples for prediction.

We show with simulated data that the fSVA approach leads to substantial improvement in predictive accuracy when unmeasured variables are correlated with biological outcomes. We also apply fSVA to multiple publicly available microarray data sets and show improvements in prediction accuracy after correcting for batch. The methods developed in this paper have been implemented in the freely available `sva` Bioconductor package.

FROZEN SURROGATE VARIABLE ANALYSIS METHODOLOGY

Removing batch effects from the training set

The first step in batch correction for prediction problems is to remove batch effects from the training set. In the training set, the biological groups are known. This setting is similar to the population genomics setting and we can use a model for gene expression data originally developed for population correction of unmeasured confounders. If there are m

Parker et al. (2014), *PeerJ*, DOI 10.7717/peerj.561

2/12

measured features and n samples in the training set, we let $X_{m \times n}$ be the matrix of feature data, where x_{ij} is the value of feature i for sample j. For convenience, we will refer to X as the expression matrix for the remainder of the paper. However, our methods can be generally applied to any set of features, including measures of protein abundance, gene expression, or DNA methylation.

We propose a linear model for the relationship between the expression levels and the outcome of interest y_j: $x_{ij} = b_0 + \sum_{k=1}^{p_1} s_k(y_j) + e_{ij}$. The $s_k(\cdot)$ are a set of basis functions parameterizing the relationship between expression and outcome. If the prediction problem is two-class, then $p_1 = 1$ and $s_1(y_j) = 1(y_j = 1)$ is an indicator function that sample j belongs to class one. In a multi-class prediction problem $k > 1$ and the $s_k(\cdot)$ may represent a factor model for each class. In matrix form this model can be written (*Leek & Storey, 2007*; *Leek & Storey, 2008*).

$$X = BS + E \qquad (1)$$

where $S_{p_1 \times n}$ is a model matrix of p_1 biological variables of interest for the n samples, $B_{m \times p_1}$ are the coefficients for these variables, and $E_{m \times n}$ is the matrix of errors.

In genomic studies, the error term E is not independent across samples (*Johnson, Li & Rabinovic, 2007*; *Leek & Storey, 2007*; *Leek & Storey, 2008*; *Walker et al., 2008*; *Friguet, Kloareg & Causeur, 2009*; *Leek et al., 2010*; *Gagnon-Bartsch & Speed, 2012*). That is, there is still correlation between rows of E after accounting for the model S. The correlation is due to unmeasured and unwanted factors such as batch. We can modify model (1) to account for these measured biological factors and unmeasured biological and non-biological factors:

$$X = BS + \Gamma G + U \qquad (2)$$

where $G_{p_2 \times n}$ is a $p_2 \times n$ random matrix, called a dependence kernel (*Leek & Storey, 2008*) that parameterizes the effect of unmeasured confounders, $\Gamma_{m \times p_2}$ is the $m \times p_2$ matrix of coefficients for G, and $U_{m \times n}$ is the $m \times n$ matrix of independent measurement errors. We previously demonstrated that such a decomposition of the variance exists under general conditions typically satisfied in population genomic experiments (*Leek & Storey, 2008*).

In the training set, the biological classes are known, so S is known and fixed. But the matrices B, Γ and G must be estimated. fSVA first performs surrogate variable analysis (SVA) on the training database in order to identify surrogates for batch effects in the training samples. The training set can be "cleaned" of batch effects by regressing the effect of the surrogate variables out of the data for each feature. Any classification algorithm can then be developed on the basis of the clean training data set.

SVA is an iterative algorithm that alternates between two steps. First SVA estimates the probabilities $\pi_{i\gamma} = \Pr(\gamma_{i\cdot} \neq \vec{0}|X, S, \hat{G}), \pi_{ib} = \Pr(b_{i\cdot} \neq \vec{0}|\gamma_{i\cdot} \neq 0, X, S, \hat{G})$ using an empirical Bayes' estimation procedure (*Leek & Storey, 2008*; *Efron, 2004*; *Storey, Akey & Kruglyak, 2005*). These probabilities are then combined to define an estimate of the probability that a

gene is associated with unmeasured confounders, but not with the group outcome

$$\pi_{iw} = \Pr(b_{i.} = \vec{0} \ \& \ \gamma_{i.} \neq \vec{0}|X, S, \hat{G})$$
$$= \Pr(b_{i.} = \vec{0}|\gamma_{i.} \neq 0, X, S, \hat{G})\Pr(\gamma_{i.} \neq \vec{0}|X, S, \hat{G})$$
$$= (1 - \pi_{ib})\pi_{i\gamma}.$$

The second step of the SVA algorithm weighs each row of the expression matrix X by the corresponding probability weight $\hat{\pi}_{iw}$ and performs a singular value decomposition of the weighted matrix. Letting $\hat{w}_{ii} = \hat{\pi}_{iw}$ the decomposition can be written $\hat{W}X = UDV^T$. After iterating between these two steps, the first p_2 weighted left singular vectors of X are used as estimates of G. An estimate of p_2 can be obtained either through permutation (*Buja & Eyuboglu, 1992*) or asymptotic (*Leek, 2011*) approaches.

Once estimates $\hat{G}$ have been obtained, it is possible to fit the regression model in Eq. (2) using standard least squares. The result are estimates for the coefficients $\hat{B}$ and $\hat{\Gamma}$. Batch effects can be removed from the training set by setting $\hat{X}^{clean} = X - \hat{\Gamma}\hat{G}$. Any standard prediction algorithm can then be applied to $\hat{X}^{clean}$ to develop a classifier based on batch-free genomic data. The result is a prediction function $f(\hat{X}_{.j}^{clean})$ that predicts the outcome variable y_j based on the clean expression matrix.

Removing batch effects from new samples

Removing batch effects from the training database is accomplished using standard population genomic SVA batch correction. But the application of classifiers to new genomic samples requires batch correction of individual samples when both the batch and outcome variables are unknown. The fSVA algorithm borrows strength from the training database to perform this batch correction.

To remove batch effects from a new sample $X_{.j'}$, it is first appended to the training data to create an augmented expression matrix $X^{j'} = [XX_{.j'}]$ where $[\cdot]$ denotes concatenation of columns. To estimate the values of G for the new sample, fSVA uses a weighted singular value decomposition, using the probability weights estimated from the training database $\hat{W}X^{j'} = U^{j'}D^{j'}V^{j'T}$. The result is an estimate $\hat{G}^{j'}$ that includes a column for the new sample. Note that only one new sample was appended. Had all the new samples been appended at once, the singular value decomposition would be highly influenced by the similarity in the new samples, rather than detecting similarities between the new sample and the database samples.

To remove batch effects from the new sample, fSVA uses the coefficients estimated from the training database $\hat{\Gamma}$ and the estimated surrogate variables: $\hat{X}^{clean\ j'} = X^{j'} - \hat{\Gamma}\hat{G}^{j'}$. If there are n training samples, then the $(n+1)$st column of $X^{clean\ j'}$ represents the new clean sample. The classifier built on the clean training data can be applied to this clean data set to classify the new sample.

FAST fSVA METHODOLOGY

fSVA requires that a new singular value decomposition be applied to the augmented expression matrix once for each new sample. Although this is somewhat computationally

intensive, in typical personalized medicine applications, sample collection and processing will be spread over a long period of time. In this setting, computational time is not of critical concern. However, for evaluating the fSVA methodology or developing new classifiers using cross-validation, it is important to be able to quickly calculate clean expression values for test samples.

We propose an approximate fSVA algorithm that greatly reduces computing time by performing a streaming singular value decomposition (*Warmuth & Kuzmin, 2007*; *Warmuth & Kuzmin, 2008*). The basic idea behind our computation speed-up is to perform the singular value decomposition once on the training data, save the left singular vectors and singular values, and use them to calculate approximate values for the right singular values in new samples.

When removing batch effects from the training data, the last step is a weighted singular value decomposition of the training expression matrix $WX = UDV^T$. After convergence, the first p_2 columns of the matrix V are the surrogate variables for the training set. Since U and V are orthonormal matrices, we can write $V^T = D^{-1}U^TWX$. The matrix $P = D^{-1}U^TW$ projects the columns of X onto the right singular vectors V^T. Pre-multiplying a set of new samples X^{new} by P results in an estimate of the singular values for the new samples: $\hat{V}^{Tnew} = P^TX^{new}$. The surrogate variable estimates for the new samples consist of the first p_2 columns of $\hat{V}^{Tnew}$. We obtain clean data for the new samples using the estimated coefficients from the training set, identical to the calculation for the exact fSVA algorithm: $\hat{X}^{clean,new} = X^{new} - \hat{\Gamma}\hat{G}^{new}$.

Estimates obtained using this approximate algorithm are not identical to those obtained using the exact fSVA algorithm. The projection matrix used in the approximation, P^T, is calculated using only the samples in the training set. However, there is only a one-sample difference between the projection calculated in the training set and the projection that would be obtained with exact fSVA. As the training set size grows, the approximation is closer and closer to the answer that would be obtained from the exact algorithm. For smaller databases, there is less computational burden in calculating the exact estimates. However, for large training sets, the computational savings can be dramatic, as described in the simulation below.

SIMULATION RESULTS

We performed a simulation to examine the benefit of fSVA in prediction problems. In order to do this, we simulated data using Eq. (2) under different distributions of each parameter. We also created discrete probability weights $\pi_{i\gamma}$ and π_{ib}, each equal to 1 to indicate batch- or outcome-affected, and 0 to indicate otherwise. We also varied the distribution of these probability weights (Table 1). We crafted these simulations to mimic scenarios with a subtle outcome and a strong batch effect, which is frequently the case in genomic data.

We also specified that both the simulated database and the simulated new samples have two batches and two outcomes. Each outcome was represented in 50% of the samples in both the database and the new samples. Similarly, each batch was represented in 50% of the database and new samples.

Table 1 Specifications for the three simulation scenarios used to show the performance of fSVA. We performed three simulations under slightly different parameterizations to show the effectiveness of fSVA in improving prediction accuracy. Parameters from Eq. (2) were simulated using the distributions specified in this table. Additionally, the percentage of features in the simulation affected by batch, outcome, or both are as indicated in this table. Results from these simulations can be found in Fig. 1.

Parameter distributions	
Scenario 1	$B \sim N(0,1)$
	$\Gamma \sim N(0,3)$
	$U \sim N(0,2)$
Scenario 2	$B \sim N(0,1)$
	$\Gamma \sim N(0,4)$
	$U \sim N(0,3)$
Scenario 3	$B \sim N(0,1)$
	$\Gamma \sim N(0,4)$
	$U \sim N(0,3)$
Affected features	
Scenario 1	50% batch-affected
	50% outcome-affected
	40% affected by both
Scenario 2	50% batch-affected
	50% outcome-affected
	40% affected by both
Scenario 3	80% batch-affected
	80% outcome-affected
	50% affected by both

In the database, we varied the amount of confounding between batch and outcome from a Pearson's correlation of 0 to a correlation of over 0.90. This mimics common database structures in publicly available repositories. Since the new samples are simulating a collection of single samples (such as new patients coming to the doctor), the correlation of batch and outcome within the new samples matrix is unimportant. To have a representative amount of new samples from each combination of batch and outcome, we found it best to simulate the new samples by leaving the batch and outcome uncorrelated. That way, each of the four test-cases of batch and outcome combinations was represented in 25% of the new samples.

We simulated 100 database samples and 100 new samples using the parameters described above. Each sample had 10,000 features. As a control, for each iteration in addition to performing fSVA correction, we performed SVA correction on the simulated database alone, and also performed prediction with no batch correction on the simulated database or new samples.

To quantify the effect that fSVA had on prediction, we performed exact fSVA as described above on the simulated database and new samples. We then performed Prediction Analysis of Microarrays (PAM), a commonly used method for classifying microarrays (*Tibshirani et al., 2002*). The PAM prediction model was built on the

SVA-corrected database, and then used to predict the outcomes on the fSVA-corrected new samples. Each simulation was repeated 100 times for robustness. We performed the simulation for a range of potential values for the correlation between the outcome we were predicting and the batch effects. The correlation quantifies how much the outcome and the batch effects overlap in the training set. When the correlation is zero, the batch effects and outcome are perfectly orthogonal. When the correlation is one, then the batch effect and outcome are the same in the training set.

We found that in general the prediction accuracy measures for different iterations of a simulation varied highly, but the ordinality remained relatively constant. Therefore to display results we randomly selected three graphs from each of the scenarios, using the `sample` function in R (Fig. 1). Each of the graphs from the 100 iterations for each scenario can be found on the author's website.

We found that fSVA improved the prediction accuracy in all of our simulations (Fig. 1). Interestingly, exact fSVA generally outperformed fast fSVA at all of the correlation levels except the highest correlation levels. However both fSVA methods out-performed our control of performing SVA on the database alone. Additionally any method of batch-correction generally outperformed no batch correction whatsoever.

When the batch and outcome were not correlated with each other, we saw ambiguous performance from using fSVA. This is not unexpected since it has been shown that in scenarios with no confounding between batch and outcome, batch has a minimal effect on prediction accuracy (*Parker & Leek, 2012*). When databases had extreme confounding between batch and outcome (above 0.85) we saw the benefits of all the batch-correction methods drop off. This is because in these situations, SVA on the database cannot differentiate batch and outcome in the database.
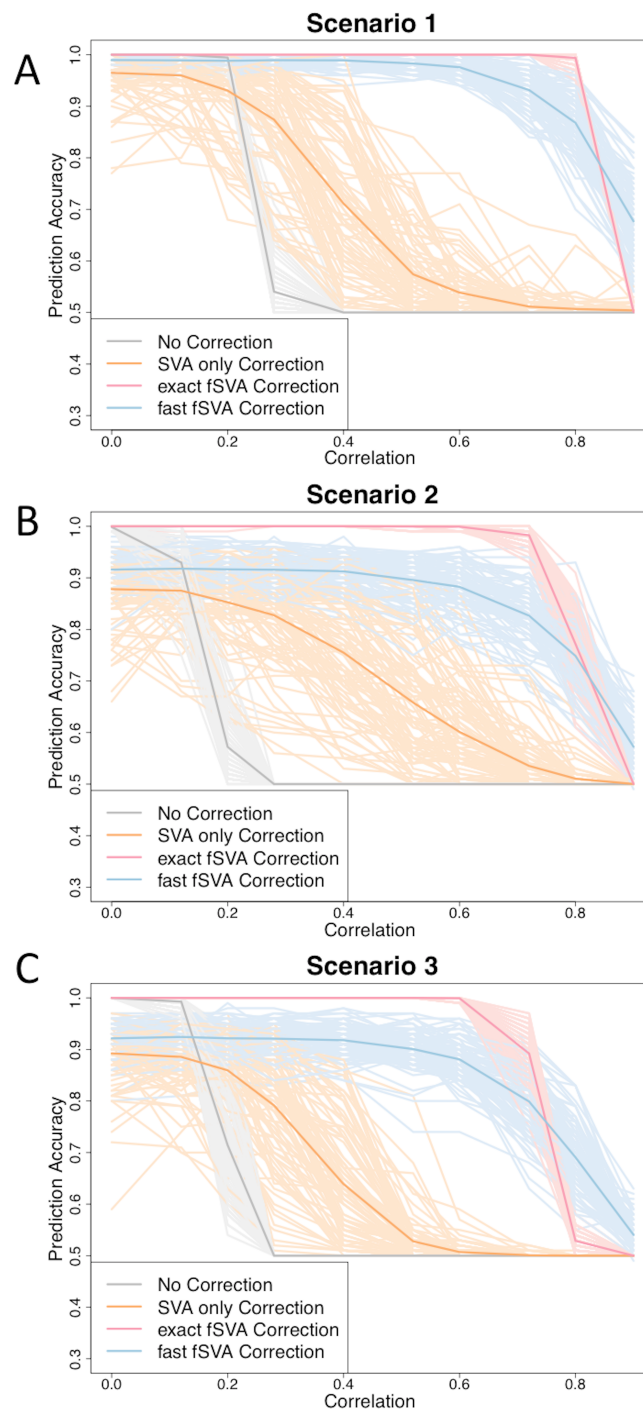
While in each of the simulations there was an accuracy cost to using fast fSVA vs. exact fSVA, the computational time savings was dramatic. In the scenario described, with 100 samples in the database and 100 new samples, the wall-clock computational time using a standard desktop computer for exact fSVA was 133.9 s, vs. just 1.3 s for fast fSVA. Using 50 samples in the database and 50 new samples, exact fSVA required 17.9 s vs. 0.4 s for fast fSVA. We encourage users to consider both the accuracy and computational times when selecting which algorithm to use for a particular data set.

RESULTS FROM MICROARRAY STUDIES

We examined the effect that fSVA had on several microarray studies, obtained from the Gene Expression Omnibus (GEO) website (*Edgar, Domrachev & Lash, 2002*). All except three of the studies were preprocessed/standardized as described previously. Three of the studies (GSE2034, GSE2603, GSE2990) were obtained from GEO and fRMA-normalized.

Each of the studies was randomly divided into equally-sized "database" and "new sample" subsets. We SVA-corrected the database subset, and then built a predictive model (PAM) on that corrected data. We then performed fSVA correction on the new samples. After performing fSVA correction, we measured the prediction accuracy of the model built on the database by calculating the number of times that the predicted outcome equaled

Figure 1 fSVA improves prediction accuracy of simulated datasets. We created simulated datasets (consisting of a database and new samples) using model (2) and tested the prediction accuracy of these using R. For each simulated data set we performed either exact fSVA correction, fast fSVA correction, SVA correction on the database only, or no correction. We performed 100 iterations on each simulation scenario described in Table 1. We performed the simulation for a range of potential values for the correlation between the outcome we were predicting and the batch effects (*x*-axis in each plot). These plots show the 100 iterations, as well as the average trend line for each of the four methods investigated.

Table 2 **fSVA improves prediction accuracy in 5 of the 9 studies examined.** The remaining 4 studies showed indeterminate results since the 95% confidence intervals overlapped zero. In order to find the prediction accuracy results, each of the studies was randomly divided into "database samples" and "new samples". Exact fSVA-correction was then performed as described above. We then built a predictive model (PAM) on the database and tested the prediction accuracy on the new samples.

Study	No correction	Improvement with fSVA
GSE10927	**0.97** (0.96, 0.97)	**0.02** (0.01, 0.02)
GSE13041	**0.61** (0.59, 0.63)	**0.07** (0.05, 0.10)
GSE13911	**0.93** (0.93, 0.94)	**0.01** (0.00, 0.01)
GSE2034	**0.51** (0.49, 0.52)	**0.03** (0.01, 0.05)
GSE2603	**0.68** (0.66, 0.70)	**−0.02** (−0.04, 0.00)
GSE2990	**0.59** (0.58, 0.61)	**−0.02** (−0.04, 0.00)
GSE4183	**0.89** (0.88, 0.91)	**−0.02** (−0.03, 0.00)
GSE6764	**0.74** (0.72, 0.76)	**0.01** (−0.01, 0.03)
GSE7696	**0.78** (0.76, 0.79)	**0.02** (0.01, 0.04)

the true outcome status, divided by the number of samples. This process was iterated 100 times for each study to obtain confidence intervals. This method is virtually identical to the simulation described above.

Results from this process can be found below (Table 2). Five of the studies showed significant improvement using fSVA. One study showed marginal improvement, with its 95% confidence interval overlapped zero. Three studies showed a cost to using fSVA, though in all three cases the 95% confidence interval for the true cost overlapped zero.

CONCLUSIONS

Batch effects have been recognized as a crucial hurdle for population genomics experiments (*Leek et al., 2010*; *Parker & Leek, 2012*). They have also been recognized as a critical hurdle in developing genomic signatures for personalized medicine (*Micheel, Nass & Omenn, 2012*). Here we have introduced the first batch correction method specifically developed for prediction problems. Our approach borrows strength from a training set to infer and remove batch effects in individual clinical samples.

We have demonstrated the power of our approach in both simulated and real gene expression microarray data. However, our approach depends on similarity between the training set and the test samples, both in terms of the genes affected and the estimated coefficients. In small training sets, these assumptions may be violated. Similarly, training sets that show near perfect correlation between batch variables and biological classes represent an extreme case that can not be directly corrected using fSVA. An interesting avenue for future research is the use of publicly available microarray data to build increasingly large training databases for batch removal. We note that when the correlation between batch effects and the outcome is high it is impossible to remove these effects entirely and this may lead to poor prediction performance (*Chikina & Sealfon, 2014*).

Therefore, it is critical to create properly designed training sets where batch effects and the outcome are not highly correlated with each other.

We have discussed here applications of fSVA to microarray data but the methodology may also be useful for other applications where genomic technologies are being applied for prediction or classification. For example, with appropriate transformations, SVA can be applied to RNA-sequencing (*Leek, 2014*) or DNA methylation (*Jaffe et al., 2012*) data. The methods we have developed here are available as part of the `sva` Bioconductor package (*Leek et al., 2012*). Code and data to reproduce this project are available at https://github.com/hilaryparker/fSVA.

ADDITIONAL INFORMATION AND DECLARATIONS

Competing Interests
The authors declare there are no competing interests.

Author Contributions
- Hilary S. Parker conceived and designed the experiments, performed the experiments, analyzed the data, contributed reagents/materials/analysis tools, wrote the paper, prepared figures and/or tables, reviewed drafts of the paper.
- Héctor Corrada Bravo contributed reagents/materials/analysis tools, wrote the paper, reviewed drafts of the paper.
- Jeffrey T. Leek conceived and designed the experiments, contributed reagents/materials/analysis tools, wrote the paper, reviewed drafts of the paper.

Data Deposition
The following information was supplied regarding the deposition of related data:
Code to reproduce our analysis: https://github.com/hilaryparker/fSVA.

Supplemental Information
Supplemental information for this article can be found online at http://dx.doi.org/10.7717/peerj.561#supplemental-information.

REFERENCES

Akey JM, Biswas S, Leek JT, Storey JD. 2007. On the design and analysis of gene expression studies in human populations. *Nature Genetics* **39**(7):807–808 DOI 10.1038/ng0707-807.

Baggerly KA, Morris JS, Coombes KR. 2004. Reproducibility of SELDI-TOF protein patterns in serum: comparing datasets from different experiments. *Bioinformatics* **20**(5):777–785 DOI 10.1093/bioinformatics/btg484.

Baggerly KA, Coombes KR, Morris JS. 2003. Bias, randomization, and ovarian proteomic data: a reply to "producers and consumers". *Cancer Informatics* **1**:9–14.

Buja A, Eyuboglu N. 1992. Remarks on parallel analysis. *Multivariate Behavioral Research* **27**(4):509–540 DOI 10.1207/s15327906mbr2704_2.

Chan IS, Ginsburg GS. 2011. Personalized medicine: progress and promise. *Annual Review of Genomics and Human Genetics* **12**:217–244 DOI 10.1146/annurev-genom-082410-101446.

Chikina MD, Sealfon SC. 2014. Increasing consistency of disease biomarker prediction across datasets. *PLoS ONE* **9**(4):e91272 DOI 10.1371/journal.pone.0091272.

Edgar R, Domrachev M, Lash AE. 2002. Gene expression omnibus: NCBI gene expression and hybridization array data repository. *Nucleic Acids Research* **30**(1):207–210 DOI 10.1093/nar/30.1.207.

Efron B. 2004. Large-scale simultaneous hypothesis testing. *Journal of the American Statistical Association* **99**(465):96–104 DOI 10.1198/016214504000000089.

Eisen MB, Spellman PT, Brown PO, Botstein D. 1998. Cluster analysis and display of genome-wide expression patterns. *Proceedings of the National Academy of Sciences of the United States of America* **95**(25):14863–14868 DOI 10.1073/pnas.95.25.14863.

Fare TL, Coffey EM, Dai H, He YD, Kessler DA, Kilian KA, Koch JE, LeProust E, Marton MJ, Meyer MR, Stoughton RB, Tokiwa GY, Wang Y. 2003. Effects of atmospheric ozone on microarray data quality. *Analytical Chemistry* **75**(17):4672–4675 DOI 10.1021/ac034241b.

Friguet C, Kloareg M, Causeur D. 2009. A factor model approach to multiple testing under dependence. *Journal of the American Statistical Association* **104**(488):1406–1415 DOI 10.1198/jasa.2009.tm08332.

Gagnon-Bartsch JA, Speed TP. 2012. Using control genes to correct for unwanted variation in microarray data. *Biostatistics* **13**(3):539–552 DOI 10.1093/biostatistics/kxr034.

Jaffe AE, Murakami P, Lee H, Leek JT, Fallin MD, Feinberg AP, Irizarry RA. 2012. Bump hunting to identify differentially methylated regions in epigenetic epidemiology studies. *International Journal of Epidemiology* **41**(1):200–209 DOI 10.1093/ije/dyr238.

Johnson WE, Li C, Rabinovic A. 2007. Adjusting batch effects in microarray expression data using empirical Bayes methods. *Biostatistics* **8**(1):118–127 DOI 10.1093/biostatistics/kxj037.

Lambert CG, Black LJ. 2012. Learning from our GWAS mistakes: from experimental design to scientific method. *Biostatistics* **13**(2):195–203 DOI 10.1093/biostatistics/kxr055.

Lander ES. 1999. Array of hope. *Nature Genetics* **21**(1 Suppl):3–4 DOI 10.1038/4427.

Leek JT. 2011. Asymptotic conditional singular value decomposition for high-dimensional genomic data. *Biometrics* **67**(2):344–352 DOI 10.1111/j.1541-0420.2010.01455.x.

Leek J. 2014. svaseq: removing batch effects and other unwanted noise from sequencing data. bioRxiv DOI 10.1101/006585.

Leek JT, Storey JD. 2007. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genetics* **3**(9):1724–1735 DOI 10.1371/journal.pgen.0030161.

Leek JT, Storey JD. 2008. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences of the United States of America* **105**(48):18718–18723 DOI 10.1073/pnas.0808709105.

Leek JT, Johnson WE, Parker HS, Jaffe AE, Storey JD. 2012. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinformatics* **28**(6):882–883 DOI 10.1093/bioinformatics/bts034.

Leek JT, Scharpf RB, Bravo HC, Simcha D, Langmead B, Johnson WE, Geman D, Baggerly KA, Irizarry RA. 2010. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics* **11**(10):733–739 DOI 10.1038/nrg2825.

Luo J, Schumacher M, Scherer A, Sanoudou D, Megherbi D, Davison T, Shi T, Tong W, Shi L, Hong H, Zhao C, Elloumi F, Shi W, Thomas R, Lin S, Tillinghast G, Liu G, Zhou Y, Herman D, Li Y, Deng Y, Fang H, Bushel P, Woods M, Zhang J. 2010. A comparison of batch effect removal methods for enhancement of prediction performance using MAQC-II microarray gene expression data. *The Pharmacogenomics Journal* **10**(4):278–291 DOI 10.1038/tpj.2010.57.

Micheel CM, Nass SJ, Omenn GS (eds.) 2012. *Evolution of translational omics: lessons learned and the path forward*. Washington, DC: National Academies Press.

Parker HS, Leek JT. 2012. The practical effect of batch on genomic prediction. *Statistical Applications in Genetics and Molecular Biology* **11**(3):1–22 DOI 10.1515/1544-6115.1766.

Scharpf RB, Ruczinski I, Carvalho B, Doan B, Chakravarti A, Irizarry RA. 2011. A multilevel model to address batch effects in copy number estimation using SNP arrays. *Biostatistics* **12**(1):33–50 DOI 10.1093/biostatistics/kxq043.

Sebastiani P, Solovieff N, Puca A, Hartley SW, Melista E, Andersen S, Dworkis DA, Wilk JB, Myers RH, Steinberg MH, Montano M, Baldwin CT, Perls TT. 2011. *Science* **333**(6041):404 DOI 10.1126/science.333.6041.404-a.

Spielman RS, Bastone LA, Burdick JT, Morley M, Ewens WJ, Cheung VG. 2007. Common genetic variants account for differences in gene expression among ethnic groups. *Nature Genetics* **39**(2):226–331 DOI 10.1038/ng1955.

Storey JD, Akey JM, Kruglyak L. 2005. Multiple locus linkage analysis of genomewide expression in yeast. *PLoS Biology* **3**(8):e267 DOI 10.1371/journal.pbio.0030267.

Tibshirani R, Hastie T, Narasimhan B, Chu G. 2002. Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences of the United States of America* **99**(10):6567–6572 DOI 10.1073/pnas.082099299.

Walker WL, Liao IH, Gilbert DL, Wong B, Pollard KS, McCulloch CE, Lit L, Sharp FR. 2008. Empirical Bayes accomodation of batch-effects in microarray data using identical replicate reference samples: application to RNA expression profiling of blood from Duchenne muscular dystrophy patients. *BMC Genomics* **9**:494 DOI 10.1186/1471-2164-9-494.

Warmuth MK, Kuzmin D. 2007. Randomized PCA algorithms with regret bounds that are logarithmic in the dimension. *Advances in Neural Information Processing Systems* **19**:1481–1488.

Warmuth MK, Kuzmin D. 2008. Randomized online PCA algorithms with regret bounds that are logarithmic in the dimension 9: 2287–2320.