



Gene expression

propeller: testing for differences in cell type proportions in single cell data

Belinda Phipson ^{1,2,3,*}, Choon Boon Sim^{4,5}, Enzo R. Porrello^{4,5,6,7}, Alex W. Hewitt^{8,9}, Joseph Powell ^{10,11} and Alicia Oshlack^{12,13,14,*}

¹Bioinformatics Division, Walter and Eliza Hall Institute of Medical Research, Melbourne, VIC 3052, Australia, ²Department of Pediatrics, University of Melbourne, Melbourne, VIC 3010, Australia, ³Department of Medical Biology, University of Melbourne, Melbourne, VIC 3010, Australia, ⁴Heart Regeneration Group, Murdoch Children's Research Institute, Royal Children's Hospital, Melbourne, VIC 3052, Australia, ⁵Melbourne Centre for Cardiovascular Genomics and Regenerative Medicine, The Royal Children's Hospital, Melbourne, VIC 3052, Australia, ⁶Department of Anatomy and Physiology, School of Biomedical Sciences, The University of Melbourne, Melbourne, VIC 3010, Australia, ⁷Novo Nordisk Foundation Center for Stem Cell Medicine, Murdoch Children's Research Institute, Royal Children's Hospital, Melbourne, VIC 3052, Australia, ⁸Menzies Institute for Medical Research, School of Medicine, University of Tasmania, Hobart, TAS, Australia, ⁹Centre for Eye Research Australia, The University of Melbourne, Melbourne, VIC, Australia, ¹⁰Garvan-Weizmann Centre for Cellular Genomics, Garvan Institute of Medical Research, Darlinghurst, NSW 2010, Australia, ¹¹UNSW Cellular Genomics Futures Institute, University of New South Wales, Kingston, NSW 2052, Australia, ¹²Bioinformatics and Computational Biology, Peter MacCallum Cancer Centre, Melbourne, VIC 3000, Australia, ¹³Sir Peter MacCallum Department of Oncology, University of Melbourne, Melbourne, VIC 3010, Australia and ¹⁴School of Biosciences, University of Melbourne, Melbourne, VIC 3010, Australia

*To whom correspondence should be addressed.

Associate Editor: Anthony Mathelier

Received on December 12, 2021; revised on August 22, 2022; editorial decision on August 22, 2022; accepted on August 23, 2022

Abstract

Motivation: Single cell RNA-Sequencing (scRNA-seq) has rapidly gained popularity over the last few years for profiling the transcriptomes of thousands to millions of single cells. This technology is now being used to analyse experiments with complex designs including biological replication. One question that can be asked from single cell experiments, which has been difficult to directly address with bulk RNA-seq data, is whether the cell type proportions are different between two or more experimental conditions. As well as gene expression changes, the relative depletion or enrichment of a particular cell type can be the functional consequence of disease or treatment. However, cell type proportion estimates from scRNA-seq data are variable and statistical methods that can correctly account for different sources of variability are needed to confidently identify statistically significant shifts in cell type composition between experimental conditions.

Results: We have developed *propeller*, a robust and flexible method that leverages biological replication to find statistically significant differences in cell type proportions between groups. Using simulated cell type proportions data, we show that *propeller* performs well under a variety of scenarios. We applied *propeller* to test for significant changes in cell type proportions related to human heart development, ageing and COVID-19 disease severity.

Availability and implementation: The *propeller* method is publicly available in the open source `speckle` R package (<https://github.com/hipsonlab/speckle>). All the analysis code for the article is available at the associated analysis website: <https://hipsonlab.github.io/propeller-paper-analysis/>. The `speckle` package, analysis scripts and datasets have been deposited at <https://doi.org/10.5281/zenodo.7009042>.

Contact: alicia.oshlack@petermac.org or hipson.b@wehi.edu.au

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Single cell RNA-sequencing (scRNA-seq) technology has led to breakthroughs in the discovery of novel cell types and enhanced our understanding of the development of complex tissues. As the technology has matured it has become relatively straightforward to profile the transcriptomes of hundreds of thousands of cells, resulting in valuable insight into the composition of tissues.

While many of the first published single cell papers focussed on defining the resident cell types in complex tissues (Bornstein *et al.*, 2018; Combes *et al.*, 2019; Liu *et al.*, 2018; Zeisel *et al.*, 2015), the field is now using this technology for complex experimental comparisons with biological replication (Bunis *et al.*, 2021; Huang *et al.*, 2021; Ren *et al.*, 2021; Sim *et al.*, 2021). Experiments with different conditions and multiple biological samples can be costly; however, substantial savings can be made by pooling cells from multiple samples. If samples are genetically diverse, they can be demultiplexed using genetic information (Huang *et al.*, 2019; Xu *et al.*, 2019). An alternative approach is to use molecular cell multiplexing protocols, such as the commercially available CellPlex from 10x Genomics. Collectively, cell multiplexing makes designing larger scRNA-seq experiments more feasible.

The first step in analysis for an scRNA-seq experiment with multiple experimental conditions and biological replicates is to identify the cell types present in each sample. However, downstream analysis requires sophisticated tools to address specific hypotheses about how a perturbation affects the biological system. Two analysis tasks are commonly performed following cell type identification in order to understand the effect of the condition. One task is to find genes that are differentially expressed between groups of samples, for every cell type observed in the experiment, similar to the analysis of bulk RNA-seq experiments (Crowell *et al.*, 2020). However, a benefit of scRNA-seq data is that we have additional information on the composition of the samples. The relative change in abundance of a cell type can be a consequence of normal development, disease, or treatment. Due to technical as well as biological sources, the cell type proportion estimates from single cell data can be highly variable. The focus of this work is to find statistically significant differences in cell type proportions between groups of samples that appropriately takes into account sample-to-sample variability.

Here, we present *propeller*, a robust and flexible linear modelling-based solution to test for differences in cell type proportions between experimental conditions. The *propeller* method leverages biological replication to obtain measures of variability of cell type proportion estimates and uses empirical Bayes to stabilize variance estimates by borrowing information across cell types. It is a flexible approach that can be applied to complex experimental designs with multiple factors. Using simulated data, we compared the performance of commonly used statistical models for testing for differences in cell type proportions in single cell data and show that *propeller* performs well across a variety of experimental scenarios. We applied *propeller* to three different single cell datasets on ageing, human heart development and COVID-19 disease severity. Our *propeller* method is publicly available in the `speckle` R package (<https://github.com/philipsonlab/speckle>).

2 The *propeller* method

Propeller is a function in the `speckle` R package that uses cell level annotation information to calculate sample level cell type proportions, followed by data transformation and statistical testing for each cell type. *Propeller* leverages biological replication to estimate the high sample-to-sample variability in cell type counts often observed in real single cell data (Fig. 1a, PBMC scRNA-seq data from 12 healthy human donors). The variability in cell type proportion estimates between samples can be large both due to technical sources, such as variation in dissociation protocols, and due to valid biological factors that contribute to variability. For example, blood cell type composition is known to change with age (Tan *et al.*, 2017). Taking into account, sample-to-sample variability when analysing differences in cell type proportions is critical as observed cell

type variances are far greater than variances estimated under a binomial or Poisson distribution, which can only account for sampling variation (Fig. 1b, PBMC scRNA-seq dataset from 12 healthy human donors).

The first step of *propeller* is to calculate the cell type proportions for each sample. *Propeller* can directly derive the counts and calculate the proportions from a Seurat or SingleCellExperiment object. This results in a matrix of proportions where the rows are the cell types, and the columns are the samples. The binomial distribution has the statistical property that proportions close to zero and one have small variance, and values close to 0.5 have large variance i.e. the variances are heteroskedastic. To overcome this, we have implemented two transformations in *propeller*: (i) arcsin square root transformation, and (ii) logit transformation. The arcsin square root transformation has the benefit that it will always produce a real value. If the logit transformation is selected an offset of 0.5 is added to the raw cell type counts matrix prior to transformation to avoid taking the log of zeroes.

Next, we test whether the transformed proportions for every cell type are significantly different between two or more experimental conditions using a linear modelling framework. If there are exactly two groups, we perform moderated *t*-tests; if there are more than two groups, we perform moderated ANOVA tests (Smyth, 2004).

These tests are moderated using an empirical Bayes framework, allowing information to be borrowed across cell types to stabilize the cell type-specific variance estimates. This is particularly effective when the number of biological replicates is small and the number of cell types is at least three (Efron and Morris, 1977), a common situation in scRNA-seq experiments. The final step in *propeller* is to calculate false discovery rates (Benjamini and Hochberg, 1995; Benjamini and Yekutieli, 2001) to account for testing across multiple cell types. The output of *propeller* consists of condition-specific proportions estimates, *P*-values and false discovery rates for every cell type observed in the experiment. The statistical details for *propeller* are described in [Supplementary Material](#).

The minimal annotation information that *propeller* requires for each cell is cluster/cell type, sample and group/condition. More complex experimental designs can be accommodated using the *propeller.test* and *propeller.anova* functions, which have the flexibility to model additional covariates of interest, such as sex or age.

3 Performance using simulated datasets

3.1 Type I error control under null simulation scenario

Although it is clear from the PBMC scRNA-seq data that cell type proportions estimates are over-dispersed (Fig. 1b), we wanted to more thoroughly evaluate the performance of *propeller* as well as other statistical methods that have commonly been used for testing differences in proportions in other fields. Using simulated cell type proportions, we compared the performance of nine different statistical models.

1. χ^2 test for differences in proportions.
2. Logistic binomial regression (special case of beta-binomial with dispersion = 0).
3. Poisson regression (special case of negative binomial with dispersion = 0).
4. *propeller* with arcsin square root transformation of proportions, denoted `propeller(asin)`.
5. *propeller* with logit transformation of proportions, denoted `propeller(logit)`.
6. Beta-binomial regression on cell type counts.
7. Negative binomial regression on cell type counts.
8. Quasi-likelihood negative binomial regression on cell type counts.
9. Centred log-ratio (CLR) transformation followed by linear regression, denoted compositional data analysis (CODA).

The first three methods do not take into account sample-to-sample variability, while the remaining six methods (4–9) do. The

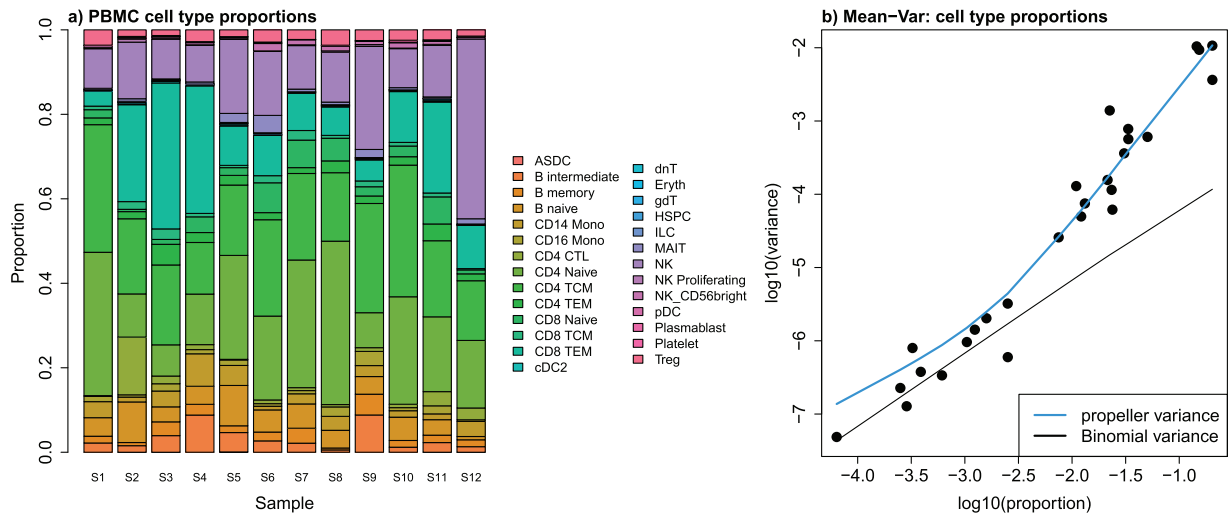


Fig. 1. Exploring heterogeneity in cell type proportions estimated from PBMC scRNA-seq data. (a) Barplot showing high levels of variability of cell type proportion estimates between 12 healthy PBMC scRNA-seq samples. (b) Mean–variance relationship for 27 cell types in 12 healthy PBMC scRNA-seq samples showing that cell type proportions are over-dispersed compared to the variance estimated under a Binomial distribution. The plot is produced using the `plotCellTypePropsMeanVar` function in the `speckle` package

quasi-likelihood approach (Method 8) is described in the Bioconductor book ‘Orchestrating single cell analysis with Bioconductor’ (<https://bioconductor.org/books/release/OSCA/>). The two variations of *propeller* model transformed proportions, while the remaining statistical tests, with the exception of the CODA method, model the cell type counts directly. Method 9 is an example from the CODA field where the cell types are modelled relative to a reference ‘cell type’. Here, the geometric mean of the cell types forms the baseline as is commonly done in microbiome data analysis (Quinn *et al.*, 2019). The log-ratio of the counts to the geometric mean is calculated and a linear model fitted with group as the explanatory variable to obtain *P*-values. The implementation of the nine methods is described in [Supplementary Material](#).

We simulated cell type counts in a hierarchical manner under a simple null scenario where the cell type proportions do not differ between two groups in order to determine whether the nine methods control the Type I error rate. We simulated five cell types with proportions that varied from rare to abundant (true proportions $\pi_i=0.01, 0.02, 0.15, 0.34, 0.45$). The sample proportions, p_{ij} , for cell type i and sample j , were generated from a Beta distribution with parameters α_i and β_i , which control how variable the proportions are. Larger values of α_i and β_i result in a more precise distribution centred around the true proportions, while smaller values result in a more diffuse prior ([Supplementary Fig. S1](#)). We set $\alpha_i=10$, and calculated the corresponding β_i for each cell type i from the following relationship derived from properties of the Beta distribution:

$$\beta_i = \alpha(1 - \pi_i) / \pi_i$$

Cell type counts x_{ij} were then sampled from a binomial distribution with parameters n_j and p_{ij} . The total number of cells, n_j , per sample j , were sampled from a negative binomial distribution with mean 5000 and dispersion 20 to simulate variation in total cell numbers per sample observed in real data. The hierarchical model is further described in [Supplementary Material](#). We simulated 10 000 datasets and counted the number of times each of the five cell types were statistically significant with *P*-value <0.05 for the nine different statistical models. We also varied the number of samples per group to determine the effect of sample size on the Type I error rate ($n = 3, 5, 10, 20$). [Figure 2a](#) shows the cell type proportions per sample observed for an example simulated dataset under these conditions.

[Supplementary Table S1](#) shows the type I error rates for the nine different methods for each of the five different cell types at a nominal *P*-value cut-off of 0.05 when the number of samples per group is

five. The most striking observation is that the statistical tests (Methods 1–3) that do not account for additional biological variability frequently find significant differences between the two groups when there are none. As expected, it is clear that methods that account for this additional variability are required and Methods 1–3 are not further explored in this analysis.

For the methods that model sample-to-sample variability none have perfect type I error rate control, although the observed rates are generally close to 0.05 (Methods 4–9). *Propeller*(asin) tends to be conservative for the most rare cell type, and permissive for more abundant cell types whereas the opposite tends to be true for the other tests, particularly for the negative binomial methods. These results show that the type I error rate differs between different cell types depending on how abundant the cell type is, and no method perfectly controls the type I error rate for both rare and abundant cell types.

[Figure 2b](#) summarizes the Type I error rates for different sample sizes. As the number of samples in each group increases, the type I error rate for all methods is closer to 0.05. For sample sizes of 10 and 20 per group, the arcsin square root transformation shows the best type I error rate control for almost all cell type abundances, however with smaller sample sizes ($n=3, 5$), the logit transform appears to better control the type I error rate. Across all sample sizes there was a trend of increased Type I error rate for the more rare cell types for *propeller*(logit), beta-binomial, negative binomial, quasi-likelihood negative binomial and the CODA method, while *propeller*(asin) tends to be conservative for the most rare cell type ($\pi=0.01$). It is not surprising that the beta-binomial model performs favourably as this method most closely resembles the distributional assumptions underlying the simulation.

3.2 Power to detect true differences in cell type proportions in simulated data

Next, we expanded the simulation to include seven cell types, three of which change proportions between the two groups by between 2 and 3-fold, while the remaining four did not ([Fig. 2c](#)). The parameters α_i and β_i of the beta distribution were estimated from real human heart single nuclei RNA-seq data ([Fig. 3a](#)) using the `estimateBetaParamsFromCounts` function available in the `speckle` package. We simulated 1000 datasets and evaluated the performance of the models by examining the proportion of simulated datasets with *P*-value <0.05 for each of the seven cell types for each of the six methods. The proportions of the three cell types that are simulated to differ between the two groups range from very rare

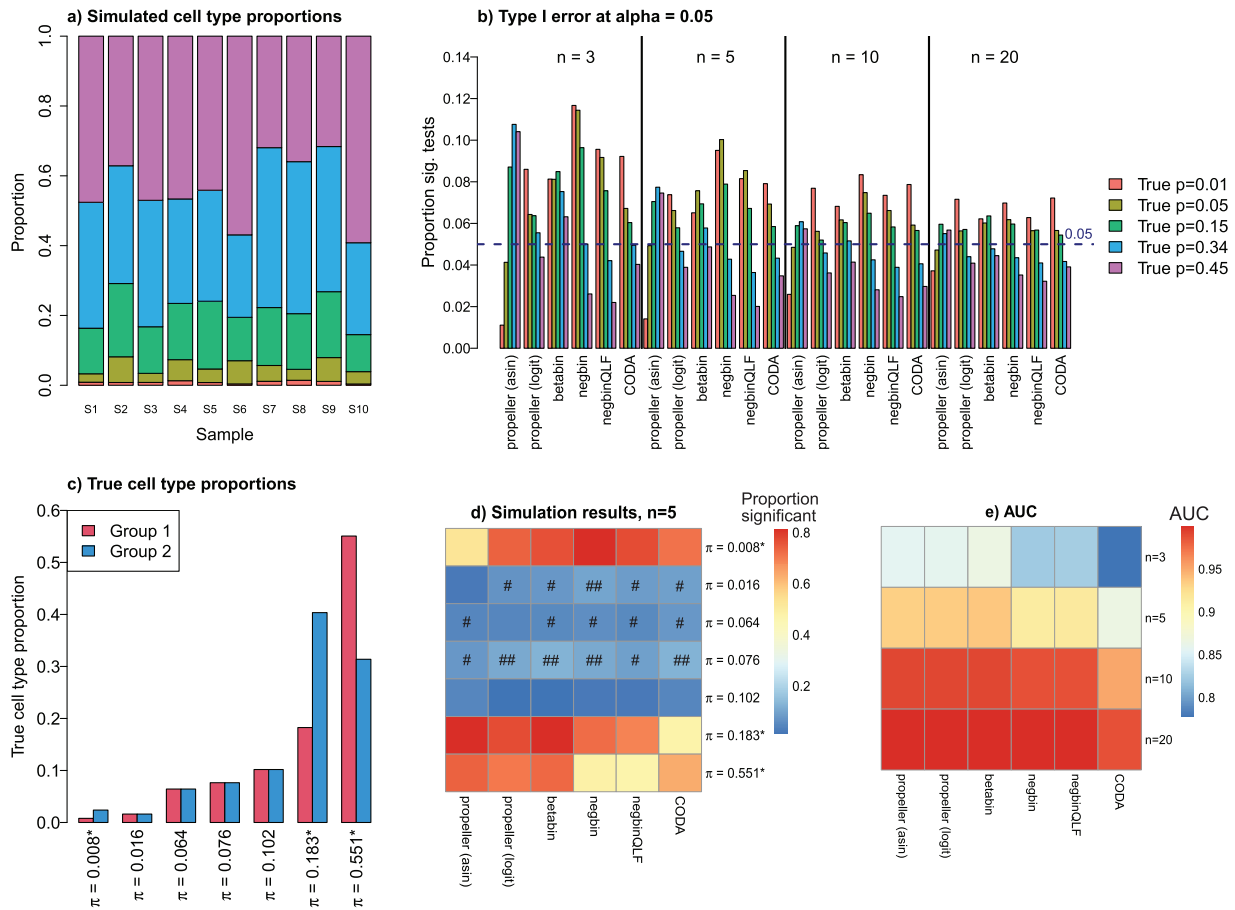


Fig. 2. Simulation results. (a) Cell type proportions for one simulated dataset with no abundance differences between Group 1 (samples S1–S5) and Group 2 (samples S6–S10). (b) Type I error rate at $\alpha = 0.05$ for sample sizes $n = 3, 5, 10, 20$ for the six methods. A total of 1000 datasets with five cell types that do not change in abundance between two groups were simulated, varying the sample size. For each of the cell types, the proportion of simulated datasets with P -value < 0.05 was calculated when testing for cell type proportion differences for each of the six models. (c) True cell type proportions for Groups 1 and 2. Three cell types that range in abundance are simulated to vary by 2–3-fold (denoted by asterisks). The remaining four cell types do not differ. (d) Heatmap showing the proportions of 1000 simulated datasets with P -value < 0.05 when testing for cell type proportion differences between two groups. True positives are denoted by an asterisk. For the true positives, a large value indicates greater power to detect significant cell type differences (proportion significant is high). For true negatives, entries without the # symbol indicates good false discovery rate control with proportion significant < 0.05 , # indicates proportion significant between 0.05 and 0.1 and ## indicates poorest control with the proportion significant > 0.1 . (e) Heatmap showing the mean AUC for each of the six methods for all sample sizes across 1000 simulated datasets. Higher AUC indicates a method has both good power to detect true positives as well as good false discovery rate control

to quite abundant (baseline proportions in Group 1 = 0.008, 0.183, 0.551). We repeated these simulations for sample sizes $n = 3, 5, 10, 20$.

At $n = 5$ samples per group, propeller(asin) detected the rare cell type difference in only 52% of the simulated datasets, while the other methods detected the rare cell type difference in 71–81% of simulated datasets (Fig. 2d). However, propeller(asin) detected the differences in the more abundant cell types in a larger percentage of the simulated datasets compared to the other methods (82% and 74% of simulated datasets). The negative binomial methods detected the difference in the most abundant cell type in $< 50\%$ of the simulated datasets. The CODA method had relatively poor performance for the more abundant cell types compared to the propeller methods and the beta-binomial model. The most consistent performing models across cell type abundances were propeller(logit) and the beta-binomial model. In terms of the cell types that did not change between the two groups, we noted that propeller(asin) generally had the best false discovery rate control and CODA had the worst. Heatmaps for sample sizes $n = 3, 10$ and 20 are shown in Supplementary Figure S2.

Figure 2e shows the mean area under the receiver operating curve (AUC) for the six methods at the four different sample sizes ($n = 3, 5, 10, 20$). As sample size increases, all methods show an improvement in performance. With at least 10 samples in each group,

all methods except CODA have an AUC above 98%. In general, propeller(asin), propeller(logit) and the Beta-binomial method have the highest AUC at each of the four sample sizes.

3.3 Extreme case: varying numbers of cell types

While the simulations above examine Type I error control and power to detect true positives with 5 and 7 cell types, respectively, we wanted to examine the performance of the methods in the extreme case when there are only 2 cell types present in the dataset, compared to when there are 20. Here, we focussed on $n = 5$ and simulating cell types with true differences between two groups.

The scenario when only two cell types are present in the data is interesting from the perspective that if one cell type changes in proportion, the other cell type will also naturally change. In this scenario, we set the Group 1 true proportions as $\pi_{1i} = 0.4, 0.6$; and Group 2 true proportions as $\pi_{2i} = 0.2, 0.8$ (Supplementary Fig. S3a). Hence, cell type 1 is halved in Group 2 compared to Group 1, and cell type 2 increases by 33.3%. In this scenario, all the methods detected the changes in the two cell types in the majority of the simulated datasets (Supplementary Fig. S3b). There was a slight decrease in power for the negative binomial methods for the more abundant cell type.

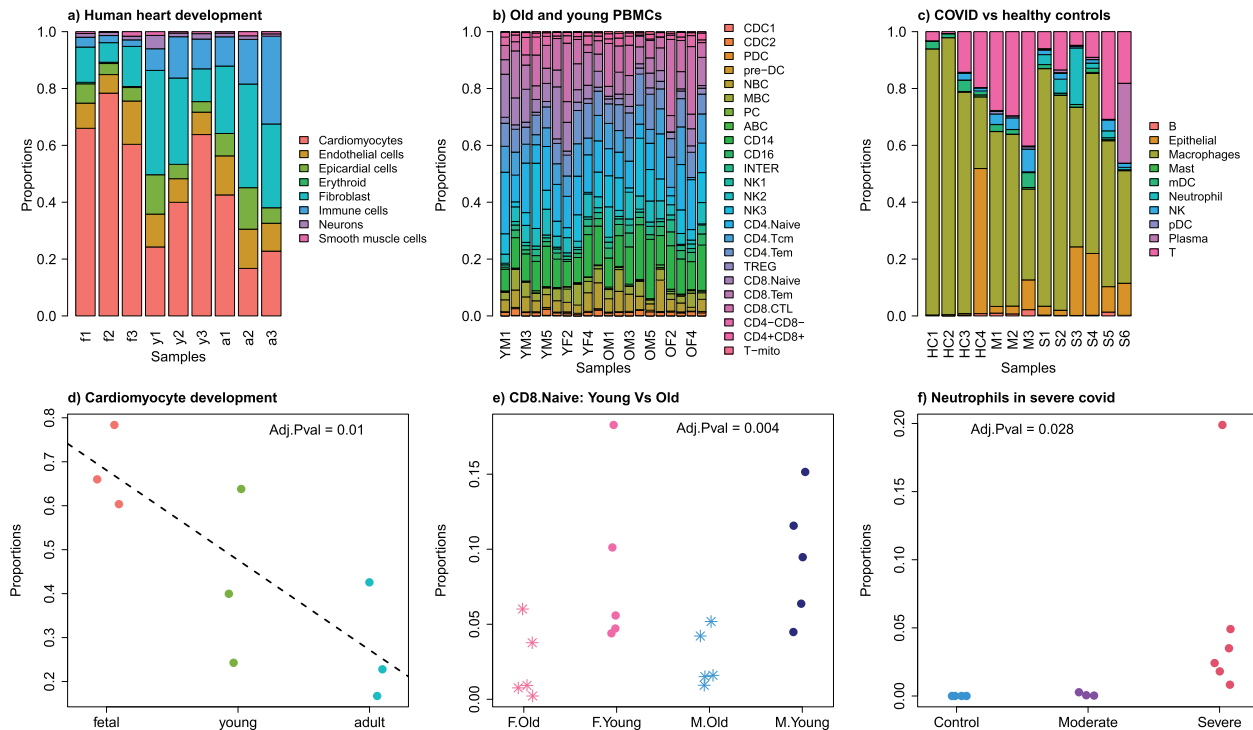


Fig. 3. Applying propeller to three scRNA-seq datasets. (a) Barplot showing cell type proportions for nine samples in a human heart development scRNA-seq dataset. f, foetal; y, young; a, adult. (b) Barplot showing cell type proportions for 20 PBMC samples that differ in terms of their age (Y/O) and sex (M/F). (c) Barplot showing cell type proportions for 13 samples in a COVID-19 study. HC, healthy control; M, moderate COVID-19 infection; S, severe COVID-19 infection. (d) Treating developmental stage as a continuous variable, the cardiomyocyte populations show a relative decline across development in human heart samples. (e) There is a statistically significant difference in the proportions of CD8 naive cells between young and old PBMC samples, taking sex into account. (f) Neutrophils are statistically significantly different between healthy control, moderate and severe COVID-19 bronchoalveolar lavage samples

For the scenario with 20 cell types, we used cell type proportion estimates from the 12 healthy human PBMC scRNA-seq dataset as our true baseline proportions. We modified 8 of the 20 cell types to be different between the two groups (Supplementary Fig. S4a). The cell types that differed in abundance between the two groups ranged from rare to abundant. The heatmap in Supplementary Figure S4b shows the proportion of significant tests across the 1000 simulated datasets for each cell type and each method. In this scenario with a larger number of cell types, the negative binomial methods have similar performance compared to the other methods. Cell types with larger log-fold changes are detected as statistically significant in the majority of simulated datasets by all methods.

Supplementary Table S2 shows the recall, precision and *F1* score for each method averaged across the 1000 simulated datasets. In this scenario, the CODA method is able to detect more true differences in cell type proportions compared to any of the other methods, at the expense of detecting the most false positives. In general, the *propeller* methods have high precision indicating that not many false discoveries are reported. The negative binomial methods perform better in this scenario, and the beta-binomial model has the second highest *F1* score with a good balance between precision and recall. Propeller(logit) has the highest precision and CODA has the highest recall.

4 Application to real single cell datasets

One important feature of *propeller* is that complex experimental designs can be modelled by using a design matrix that takes account of multiple factors. In order to demonstrate the types of experimental designs that can be accommodated, we applied propeller(logit) to three different scRNA-seq datasets that varied in terms of the experimental design and the number of samples and cell types in each dataset:

1. Nine human heart biopsy samples across development (foetal, young and adult), with eight broad cell types annotated (Sim

et al., 2021). We modelled development as a continuous variable and sex as a categorical variable.

2. A total of 20 PBMC samples across young and old male and female samples with 24 cell types annotated (Huang *et al.*, 2021). We modelled age and sex as categorical variables.
3. A total of 13 bronchoalveolar lavage fluid immune cell samples across three groups (healthy controls, moderate and severe COVID-19 infection) with 10 cell types annotated (Liao *et al.*, 2020). We modelled disease status as a categorical variable and performed an ANOVA to find cell type differences between the three groups.

Figure 3(a–c) shows the cell type proportion estimates for each sample for the three different datasets. The cell type proportions are highly variable between individuals across all datasets. Across healthy human heart development, we detected significant changes in the abundances of immune, erythroid, cardiomyocyte and fibroblast cells (Supplementary Fig. S5). In the original analysis, propeller(logit) was applied as an ANOVA test, ignoring sex. While the conclusions are not markedly different, the order of significant cell types has changed with immune cells the most significant cell type when modelling development as a continuous variable. The immune and erythroid cell type changes across development form a type of positive control and it is encouraging that they are the most significant cell types. As noted in the initial paper (Sim *et al.*, 2021), immune cells increase throughout development, as would be expected as the foetus has not been exposed to many pathogens, while an adult would have a larger and more diverse repertoire of immune cells. With the erythroid cells, only foetal red blood cells are nucleated, and hence they are captured with the nuclei protocol in foetal samples and absent in young and adult samples. An interesting finding in this study was that the relative abundance of cardiomyocytes declines with age (Fig. 3d), while fibroblasts increase across development (Supplementary Fig. S6).

For the ageing PBMC dataset, we detected statistically significant differences in CD8 naive and CD16 cells between young and old samples, while controlling for sex, at a false discovery rate threshold of 0.05. CD8 naive cells were significantly enriched in the young samples (Fig. 3e), and CD16 cells were significantly depleted in the young samples compared to old (Supplementary Fig. S6). While the CD8 naive result was reported in the initial paper, we detected a significant change in abundance of CD16 cells between young and old samples that was not reported in the original analysis (Huang *et al.*, 2021).

For the COVID-19 dataset, we found four cell types to have significant changes in abundance between the three groups when we applied propeller(logit) (Supplementary Fig. S7). We found that neutrophils were the most significantly different between healthy controls and moderate and severe bronchoalveolar lavage samples from COVID-19 patients (Fig. 3f), and this was not reported as statistically significant in the original analysis (Liao *et al.*, 2020). Plasma, pDC and NK cells also showed significant changes in abundance. Upon closer inspection, it appeared that the significant result for Plasma was driven by one sample in the severe COVID-19 group (Supplementary Fig. S7). When we re-analysed the data with propeller(asin), this cell type was no longer significant, while neutrophils, pDC and NK cells were still statistically significant. This indicates that propeller(logit) may be sensitive to outlier samples and suggests that propeller(asin) is a more robust method to use when outliers are present in the data. Compared to the results from the original analysis, we detected two additional cell types, neutrophils and NK cells, that significantly changed in abundance between healthy controls, moderate and severe COVID-19 patients.

5 Discussion

In this article, we present *propeller*, a new method for testing for differences in cell type proportions from single cell data. It takes account of sample-to-sample variability, which is large due to both technical and biological sources. The *propeller* function itself interoperates with Seurat and SingleCellExperiment class objects, and the *propeller.ttest* and *propeller.anova* functions have the ability to model complex experimental designs. In order to work specifically with the features of single cell data, which often have extreme cell type proportions, we have implemented *propeller* with two different transformations: the arcsin square root transformation and the logit transformation. Through simulation studies, we found propeller(logit) has superior performance in terms of power to detect changes in cell type proportions, as well as good false discovery rate control. Through analysis of real datasets, we found that propeller(logit) may be sensitive to outlier samples, while propeller(asin) is not, which suggests that propeller(asin) is a good alternative in this scenario. A recent comparison of statistical methods for performing cell type composition analysis of single cell data found that propeller(asin) and Dirichlet regression had the best performance across a variety of scenarios (Simmons, 2022). The *propeller* methods have the ability to handle zeros and ones in the data, which are not uncommon in cell type proportion estimates from single cell data. Zero values need to be carefully dealt with when using CODA methods. For the simulation studies, we replaced zeroes with 0.5 prior to CLR transformation. Another factor to consider when using a CODA framework is the choice of reference cell type, and all results need to be interpreted relative to the reference cell type, which can make interpretation of the output more challenging.

In our simulation studies, we explored the effect of the number of cell types on the performance of the methods. For datasets with fewer cell types, the negative binomial methods and the CODA method show decreased performance compared to beta-binomial and propeller methods. As the number of cell types increases to 20, the performance of negative binomial and CODA methods improve to be comparable to the other methods. We also explored the effect of sample size and baseline abundance of the cell type on the performance of the methods. For small sample sizes the differences

between the methods are more pronounced, with propeller(logit) and beta-binomial showing the best overall performance. As the sample size increases beyond 10 samples per group, all methods show good power and false discovery rate control, with the exception of the CODA method, which has increased false discovery rates for all cell types with increasing sample size. We also found that at smaller sample sizes, propeller(asin) had less power to detect the difference in the most rare cell type, while the negative binomial methods had less power to detect differences in the most abundant cell types.

We applied propeller to the analysis of three different single cell datasets that differed in terms of tissue, number of cell types, sample size and experimental conditions. We found significant biological differences in abundance, including some cell types that had not been previously reported in three different studies: across healthy human heart development, comparing blood from young and old patients, and in lung fluid from individuals with severe covid versus moderate and healthy controls. All our analysis is available via a *workflowr* (Blischak *et al.*, 2019) website (<https://phipsonlab.github.io/propeller-paper-analysis/>), with the original source code available on github (<https://github.com/phipsonlab/propeller-paper-analysis/tree/master>). The propeller methods are available in the speckle R package (<https://github.com/phipsonlab/speckle>).

Acknowledgements

We thank Xinyi Jin for preparing the speckle package for submission to Bioconductor. We thank Dr Jovana Maksimovic for suggesting the COVID-19 dataset as an application dataset, and Prof. Gordon Smyth for discussion regarding multiple testing adjustment under dependence of test statistics.

Funding

This work was supported by a National Health and Medical Research Council Investigator grant [GNT1175653 to B.P.]; a National Health and Medical Research Council Investigator grant [GNT1196256 to A.O.]; a National Health and Medical Research Council Investigator grant [GNT1175781 to J.P.]; and a National Health and Medical Research Council Investigator grant [GNT2008376 to E.R.P.]. The work was supported by National Health and Medical Research Council Grant [GNT1187748]. The snRNA-seq dataset for heart was funded by Royal Children's Hospital Foundation and National Health and Medical Research Council Project grant [GNT1160257]. The Novo Nordisk Foundation Center for Stem Cell Medicine is supported by Novo Nordisk Foundation grants [NNF21CC0073729 to E.R.P.]. MCRI is supported by the Victorian Government's Operational Infrastructure Support Program.

Conflict of Interest: none declared.

Data availability

The data underlying this article are available in Zenodo at <https://doi.org/10.5281/zenodo.7009042>.

References

- Benjamini, Y. and Hochberg, Y. (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Series B Stat. Methodol.*, 57, 289–300.
- Benjamini, Y. and Yekutieli, D. (2001) The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.*, 29, 1165–1188.
- Blischak, J.D. *et al.* (2019) Creating and sharing reproducible research code the workflow way. *F1000Res.*, 8, 1749.
- Bornstein, C. (2018) Single-cell mapping of the thymic stroma identifies IL-25-producing tuft epithelial cells. *Nature*, 559, 622–626.
- Bunis, D.G. *et al.* (2021) Single-cell mapping of progressive fetal-to-adult transition in human naive T cells. *Cell Rep.*, 34, 108573.

- Combes, A.N. *et al.* (2019) Single cell analysis of the developing mouse kidney provides deeper insight into marker gene expression and ligand-receptor crosstalk. *Development*, **146**, dev178673.
- Crowell, H.L. *et al.* (2020) Muscat detects subpopulation-specific state transitions from multi-sample multi-condition single-cell transcriptomics data. *Nat. Commun.*, **11**, 6077.
- Efron, B. and Morris, C. (1977) Stein's paradox in statistics. *Sci. Am.*, **236**, 119–127.
- Huang, Y. *et al.* (2019) Vireo: Bayesian demultiplexing of pooled single-cell RNA-seq data without genotype reference. *Genome Biol.*, **20**, 273.
- Huang, Z. *et al.* (2021) Effects of sex and aging on the immune cell landscape as assessed by single-cell transcriptomic analysis. *Proc. Natl. Acad. Sci. USA*, **118**, e2023216118.
- Liao, M. *et al.* (2020) Single-cell landscape of bronchoalveolar immune cells in patients with COVID-19. *Nat. Med.*, **26**, 842–844.
- Liu, Y. *et al.* (2018) Single-cell RNA-seq reveals the diversity of trophoblast subtypes and patterns of differentiation in the human placenta. *Cell Res.*, **28**, 819–832.
- Quinn, T.P. *et al.* (2019) A field guide for the compositional analysis of any-omics data. *GigaScience*, **8**, giz107.
- Ren, X. *et al.* (2021) COVID-19 immune features revealed by a large-scale single-cell transcriptome atlas. *Cell*, **184**, 1895–1913.e19.
- Sim, C.B. *et al.* (2021) Sex-specific control of human heart maturation by the progesterone receptor. *Circulation*, **143**, 1614–1628.
- Simmons, S. (2022) Cell type composition analysis: comparison of statistical methods. bioRxiv (p. 2022.02.04.479123). <https://doi.org/10.1101/2022.02.04.479123>.
- Smyth, G.K. (2004) Linear models and empirical bayes methods for assessing differential expression in microarray experiments. *Stat. Appl. Genet. Mol. Biol.*, **3**, Article3.
- Tan, Q. *et al.* (2017) Handling blood cell composition in epigenetic studies on ageing [review of handling blood cell composition in epigenetic studies on ageing]. *Int. J. Epidemiol.*, **46**, 1717–1718.
- Xu, J. *et al.* (2019) Genotype-free demultiplexing of pooled single-cell RNA-seq. *Genome Biol.*, **20**, 290.
- Zeisel, A. *et al.* (2015) Brain structure. Cell types in the mouse cortex and hippocampus revealed by single-cell RNA-seq. *Science*, **347**, 1138–1142.