The sum of two halves may be different from the whole. Effects of splitting sequencing samples across lanes.

Eleanor C. Williams^{1,*}, Ruben Chazarra-Gil^{1,*}, Arash Shahsavari^{1,*}, and Irina Mohorianu^{1,†}

¹Wellcome - MRC Cambridge Stem Cell Institute, Jeffrey Cheah Biomedical Centre, University of Cambridge, CB2 0AW, UK

*Joint first authors

†Corresponding author: I Mohorianu, iim22@cam.ac.uk

Abstract

Over the past two decades, the advances in high throughput sequencing (HTS) enabled the characterisation of biological processes at an unprecedented level of detail; as a result the vast majority of hypotheses in molecular biology rely on analyses of HTS data. However, achieving increased robustness and reproducibility of results remains one of the main challenges across analyses. Although variability in results may be introduced at various stages, such as alignment, summarisation or detection of differences in expression, one source of variability has been systematically omitted: the consequences of choices that influence the sequencing design which propagate through analyses and introduce an additional layer of technical variation.

In this study, we illustrate qualitative and quantitative differences in results arising from the splitting of samples across lanes, on bulk and single cell sequencing outputs. For bulk mRNAseq data, we focus on differential expression and enrichment analyses; for bulk ChIPseq data, we investigate the effect on peak calling, and the peaks' properties. At single cell level, we concentrate on the identification of cell subpopulations (cells clustered based on their expression profiles). We rely on the identity of markers used for assigning cell identities; both smartSeq and 10x data are presented.

We conclude that the observed reduction in the number of unique sequenced fragments reduces the level of detail on which the different prediction approaches depend. Further, the sequencing stochasticity adds in a weighting bias corroborated with variable sequencing depths.

Keywords: mRNAseq, ChIPseq, smartSeq, 10x, sample splitting, differential expression, enrichment analysis, cell type calling

Introduction

The recent developments and improvements in high-throughput sequencing (HTS) technologies facilitated increasingly complex transcriptome/genome-wide analyses [1], enhancing both the qualitative annotation of genomes [2, 3, 4] and their quantitative, functional characterization through differential expression studies [5, 6]. The diversification of methods specialised to a wide range of perspectives on DNA/RNA biology [7] was complemented by studies at single cell level [8]. While advances were observed across all aspects of the sequencing workflow [9], the amount of data produced increased most rapidly. This created another challenge: producing robust and reproducible results and simultaneously keeping up with the increasing intricacy of experiments [10].

The variability of sequencing output, which propagates through to quantification and other downstream exploration, poses one of the main challenges in bioinformatics analyses, since it implies the disentangling of relevant from irrelevant sources of variation. While the biologically relevant quantities are context dependent [11], an essential distinction exists between variability due to biological processes and variability due to measurement error or inaccuracy [12, 13]. The former is generally specific and well defined in relation to a condition; even when it is perturbed by noise, an underlying pattern of expression may emerge [14]. Technical variability encompasses measurement error [15], sequencing bias [16], and variability due to missing data [17]. For the latter the assessment of technical variation can be hindered by the lack of a ground truth.

Several studies proposed approaches to identify and characterise the sources of variability in HTS experiments, focusing on several aspects of signal distribution, which can affect the accuracy of the downstream analyses and interpretations, and jeopardise the reproducibility of results [11, 18]. These included both the analysis of noise [14, 19] and the downstream components of the analyses such as batch/background effect [20], alignment approaches [21], processing pipelines [5], normalisation methods [22] and differential expression thresholds [23, 24]. To model the intrinsic biological variability, the number of replicates in the context of experimental design was optimised using power calculations [11, 25], designed to provide a robust estimation of differential expression. These approaches rely on simulations on the number of expressed genes, on mean-dispersion estimates and dropouts after applying frequency and outlier filtering; traditional approaches do not take into account elements of sequencing design including across-lane sample splitting. In general, the impact of library construction and flow cell and lane characteristics on downstream analysis has not been studied in detail.

Sources of technical variability for RNAseq experiments span from the combinatorial numbers of highly variable isoforms to the handling of ambiguous or multi-mapped reads [1]. For ChIP datasets, the ability to address specific biological questions can be significantly impacted by antibody efficiency and specificity [26] as peak distributions are a direct consequence of affinity, over-crosslinking, DNA fragmentation and PCR amplification; for such samples, users are faced with a trade-off between number of usable reads (sensitivity of peak detection) and proportion of false positives derived from multi-mapped reads [27]. Low quality replicates can also generate bottlenecks when used in conjunction with good samples, as true peaks missing from poor quality replicates may be marked as non-

reproducible, thus creating false negatives [26]. Single-cell experiments share some of the drawbacks of bulk ones; in addition, the exponential increase in the number of cells profiled per study, coupled with the shallower sequencing depth, redefined some of the known difficulties, such as the characterisation of noise [28, 19].

In this paper, we investigate the effect of across-lane sample splitting, at the sequencing stage, on downstream analyses for bulk and single cell data. We study the differences between ground truth and samples on which we simulate across-lane sample splitting, with various parameters controlling the number of splits and proportion of reads produced on each lane. We focus on standard analyses i.e. identification of differentially expressed (DE) genes or ChIP peaks for bulk analyses; for single cell analyses, we concentrate on the allocation of cells to clusters (viewed as proxies for cell types), and comment on the observed variability in biological interpretations.

Materials

We illustrate the split effects on bulk and single-cell mRNA data, and bulk ChIPseq data. For the former (bulk mRNAseq) we used the Yang et al 2019 dataset [29], focusing on the 0h and 12h samples (GSE117896, comprising SRR7624365, SRR7624366 (biological replicates for 0h), SRR7624371 and SRR7624372 (biological replicates for 12h)). The bulk ChIPseq analysis was performed on H3K4me3 and H3K27ac samples, using 0h and 12h samples for each (SRR7624381, SRR7624384, SRR7624389, SRR7624392).

To exemplify the effect on plate-based scRNA-seq platforms (smartSeq) we use the Cuomo et al 2020 dataset [30]. We selected data from 6 donors, on 4 time points. On this input six experimental study cases (Table 1) were designed to illustrate the effect of the different covariates i.e. the donor, the specific time-point and cell types (resulting clusters) identities.

	Study case	N cells	Donor	N cells (donor)	Time point	N cells (time point)	Cluster	N cells (cluster)
1	Case 1	105	hayt	105	day2	105	0	105
2	Case 2	106	pahc	106	day3	106	1 4	66 40
3	Case 3	168	melw qunz	94 74	day0	168	3	168
4	Case 5	168	hayt	168	day1 day3	61 107	9 1	61 107
5	Case 6	95	melw vils	47 48	day1	95	5 6	45 50
6	Case 8	217	melw naah	95 122	day0 day3	95 122	3 2	95 122

Table 1: Overview of the study cases built on the scRNA-seq Smart-seq data (Cuomo et al 2020), illustrating the combinations of the different covariates. Study case 4 (same donor, different time-point and same cluster) and case 7 (different donor, different time-point, and same cluster) could not be generated due to the data structure.

To investigate the effects of lane-splitting on 10x Genomics scRNA-seq data, we used an in-vivo dataset of human hematopoietic stem and progenitor cells from spleen, bone marrow, and peripheral blood of two donors (we refer to

donor SAMEA6646089 as P1 and donor SAMEA6646090 as P2) [31]; the data is available via BioStudies accession S-SUBS4.

Methods

[splitting strategy] The splitting strategy is consistent across all datasets. The splitting is performed per sample. For each dataset, on the ground truth (GT), i.e. the original sample, and on the simulated, split (S) samples, the same pipelines with identical, default parameters are applied, to ensure an unbiased comparison between results obtained on the GT and S samples, respectively.

Let k be the number of splits for a sample, and the number of iterations n , are reserved as user-defined hyper-parameters. The steps for generating an S sample are: [1] subsampling reads without replacement from the GT sample to $\frac{100}{k}\%$ of GT in n iterations; [2] subsequently the k subsamples are concatenated. For bulk mRNAseq and ChIPseq samples, we assessed $k = 2$ and $k = 3$, $n = 10$. In Discussion, we present simulated samples for $k = 2, 3, 4, 5$ for one GT sample (bulk mRNAseq, 0h rep 1), $n = 10$; we also analysed simulated samples for $k = 2$ with variable split proportions from 55 - 45 to 90 - 10. For the bulk ChIPseq data we assessed $k = 2$ and $k = 3$, $n = 10$. For the smartSeq data and 10x data we generated $n = 3$ S samples for $k = 2$, for each of the study cases (i.e. each subset of samples), respectively, using the seqtk toolkit (<https://github.com/lh3/seqtk>).

[bulk mRNAseq] The raw samples (no QC-based filters applied) were aligned to the *M. musculus* genome [Ensembl 98.38] [32] using STAR 2.7.0a (paired-end mode) [33]. Next, the expression was quantified using featureCounts 2.0.0 [34]. To assess the stability between S and GT samples, abundance density plots and MA plots were produced. We used noisyR [14] to perform noise analysis on the GT and S samples, and further analysed the PCC distribution for binned abundances. For each dataset, expression levels were normalised using quantile normalisation [35] and DE genes between 0h and 12h were identified using edgeR [36]. The DE call was based on $|\log_2\text{FC}| > 0.5$ and adjusted p-value < 0.05 (using Benjamini-Hochberg multiple testing correction). Enrichment analysis on the resulting DE genes was performed using the gprofiler2 package [37] on GO, KEGG, reactome and transcription factor terms, with the background set as the set of expressed genes with abundance > 0 in at least one sample.

[bulk ChIPseq] The GT and S samples were aligned to the *M. musculus* genome using STAR 2.7.0a (single-end mode using default parameters) [33]. Narrow peaks were called using macs2 2.2.7.1 [38]. Peaks were matched between samples if the midpoint of the peak in sample 1 is within the boundaries of the peak in sample 2 and vice versa. Across sets of samples, amplitudes were normalised using quantile normalisation [35] and differentially methylated peaks between 0h and 12h were identified using edgeR [36]. Peaks were called DE if $|\log_2\text{FC}| > 0.5$ and adjusted p-value < 0.05 (using Benjamini-Hochberg multiple testing correction).

[sc smartSeq] Six experimental study cases were designed based on the Cuomo dataset [30], considering the donor,

time-point and cluster cell identities (Table 1). The corresponding samples were standardised in read length using Trim Galore (0.4.1) by: trimming 10bp from the 5' end (to reduce the effect of sequencing bias) and 40 bp from the 3' end (to address high adapter sequence content); the length of the resulting reads was 75 nts. The GT and S samples were aligned to the *H. sapiens* genome (GRCh38.p13) [32] using STAR (2.7.0a) [33] in paired-end mode. The gene counts were summarised in a matrix using featureCounts 2.0.0 [34]. We applied fastQC [39] to obtain read quality metrics and multiQC [40] to aggregate QC results from the reads, alignment and quantification. Seurat objects [41] were created considering features expressed in > 3 cells, and cells with > 50 features. The analysis pipeline comprises: i) normalization of expression levels (SCTransform [42]), ii) computation of PCA and UMAP embeddings (RunPCA, RunUMAP), iii) Neighborhood graph computation (FindNeighbors), iv) Clustering (FindClusters), v) Differential Expression (FindAllMarkers). We considered as DE features those with $\log_2FC > 0.5$ (i.e. the positive markers) and an adjusted p-value < 0.05 (using Bonferroni multiple testing correction). To assess the similarity of the partitioning in the GT and S samples, we calculated Jaccard similarity indices (JSI, [43]) on the cluster-specific sets of DE features (Fig 2D); restricting the JSI to the smaller set of markers.

[sc 10x Genomics] The GT and S fastq files were aligned to the 10x *H. sapiens* GRCh38 v3 reference transcriptome; the protein-coding genes were quantified using 10x Cellranger v3.1.0 [44]. The processing and analysis was performed individually for each donor, for the GT and S samples: after inspection of distributions of UMIs, number of detected genes, proportion of UMIs from mitochondrial genes (MT) and ribosomal protein-coding genes (RP), only cells with $> 1,000$ unique genes, $< 10\%$ MT and $> 20\%$ RP were retained for downstream analysis; the MT and RP genes were subsequently discarded from the count matrix. Following normalization using SCTransform [42], the 3,000 most abundant genes, accounting for 60% – 85% of UMIs in cells across the data, were identified and used for the calculation of PCA. A UMAP dimensionality reduction was calculated using the 30 first PCs (after inspection of an elbow plot of PC variance); the UMAP was subsequently used to assess the extent of potential batch effects originating from the tissue origin of cells, raw and normalized sequencing depths, MT% and RP%.

A 20-nearest neighbour graph was computed on the first 30 PCs of the data; the cells were clustered using SLM community detection [45] on the NN graph. To assess the clustering similarity, element centric clustering comparison [46] was employed on the set of common barcodes between GT and S using the ClustAssess R package [47]. Cluster markers were identified using the ROC test in Seurat v3.1.4 [41]; only genes with $|FC| > 2$ were considered. The top 25 markers per cluster, ranked by discriminative power, were subsequently used to calculate the per-cell JSI between cluster markers; the JSI when at least one of the sets is empty was set, by default, to 0.

Results

Effects of across-lane splitting on bulk data

[bulk mRNAseq] For the bulk mRNAseq study case [29], we focus on the effect of across-lane sample splitting on expression quantification, DE calling and enrichment analysis (summarised in Fig 1A-F). First, we assess the differences in quantification corresponding to the same time-point/biological replicate, 0h (rep 1), on GT vs $k = 2$ S sample S_2 (1A) and $k = 2$ vs $k = 3$ samples S_2 and S_3 (1B). We observe standard MA funneling shapes, with a variation in excess of 1.5FC for low abundances for both comparisons. In addition, for low to medium abundances, the $S_2 - S_3$ comparison yields a wider distribution of $\log_2$ FCs underlining the technical variation (noise) that is introduced (Fig 1B). Also noted is a slight shift of FCs (Fig 1A) driven by the stochastic redistribution of reads i.e. lack of signal for some low abundance genes, and over-expression of medium-high abundance genes. Similar conclusions are presented in Supplementary Fig 1A and 1B illustrating GT vs $k = 3$ samples $S1_3$ and two simulations of $k = 2$ $S1_2$ and $S2_3$, respectively.

Next, we focus on the DE call between the 0h vs 12h replicates, determined using standard pipelines and parameters. We illustrate the properties of the DE sets called on GT vs $k = 2$ (Fig 1C) and $k = 2$ vs $k = 3$ (Fig 1D), respectively; overall the results converge i.e. all genes are located in the proximity of the diagonal, however the variation is larger for the $k = 2$ vs $k = 3$ simulation $S1_2$ vs $S1_3$ (1D). Specifically, we see 10,810 genes (31.2% of expressed genes) with $> 0.5 \log_2$ FC absolute difference for $S1_2$ vs $S1_3$, compared to 7737 (22.3%) between GT and $S1_2$. The colour gradient is proportional to the average abundance of genes and underlines that medium abundance genes are mostly affected by the variability in fold-change amplitude and identity. We also assessed the stability of DE signal, for the GT and two sets of simulations for $k = 2$ and $k = 3$, summarised as an upset plot (the $x/5$ intersections are represented, with $x \in \{5, 4, 3, 2, 1\}$). While the majority of genes are detected across all comparisons (746 out of 920 for DE in GT and 1318 called DE in any GT or S sample), we notice comparison-specific genes (up to 81 out of 1012 called DE for $S2_3$). In particular, this analysis revealed only few "false negatives" (4 genes DE only in GT samples) in contrast to a larger number of "false positives" (between 48 and 81 from each S sample). 1 gene (out of 1318 called DE in any GT or S sample) was identified across all S samples but not in the GT. To link back the variability of signal to the biological interpretation of the results, we compare the sets of enriched terms, corresponding to the various DE predictions (Fig 1F). Similarly to the DE summary (Fig 1E), we observe a 68.3% consistency in results i.e. significant terms shared between all comparisons (out of 101 identified in total for GT, 141 identified for at least one comparison). However, "false positives" entries are present (24 in total, between 2 and 11 from each S sample) identified only for one set of S samples. These corroborated results underline a mixture of consistent patterns (on S samples) and random variation, difficult to predict or mitigate. This suggests lane-splitting introduces non biologically-robust results.

[bulk ChIPseq] The effects of across-lane sample splitting are also observed on ChIPseq data; for this case-study, we will focus on the changes in peak amplitude and length; H3K4me3 (Fig 1G-K and Supplementary Fig 2) and H3K27ac (Supplementary Fig 3) data will be exemplified [29]. The variability on peak calling for H3K4me3 data is assessed using MA (Fig 1G,H) and cross plots (Fig 1I,J). Similarly as for the mRNAseq data, we observe a funneling behaviour when GT and S samples are compared (Fig 1G and H), a wider variation in FC for S_{12} vs S_{13} (Fig 1H), and a over-amplification of peak abundances in the simulated samples (Fig 1G); additional supporting results are presented in Supplementary Fig 2. These fluctuations propagate on DE calls (performed using standard pipelines, on the 0h and 12h replicates), and are summarised in Fig 1I and 1J. Despite the overall convergence, for the H3K4me3 data we notice a wider variability than for the mRNAseq data, with clear "false negative" peaks in the simulated data i.e. medium abundance peaks identified as DE in the GT comparison, with a $\log_2FC > 2$, and not detected as DE in the S samples (borderline vanishing peaks); this behaviour observed also for H3K27ac data (Supplementary Fig 3C,D).

To further investigate the reduced robustness for the DE call, we focused on other peak properties such as peak length (defined as stop - start coordinates from the macs2 narrowPeak output); in Fig 1K we illustrate the distribution of ratios of peak lengths when GT and S samples are compared. We notice a global, systematic shift towards shorter peaks in simulated samples, the effect is more pronounced for $k = 3$, and a stability of behaviour across simulations (blue and purple distributions). The decrease in peak length can only be observed for common peaks. The reduction in number of called peaks is illustrated in Supplementary Fig 2E and 2F. For the 0h sample (SRR7624381), we see 10846 fewer peaks called for S_{12} samples (mean across 10 iterations) compared to GT and 14885 fewer for S_{13} samples compared to GT, yielding percentage decreases averaging 12.7% and 18.4% respectively.

Simulated across-lane sampling splitting has multiple effects on the peaks called in ChIPseq data, from the number of peaks to their amplitude and length; this technical, stochastic variation may have an impact on the interpretation of results.

Figure 1: **Comparison of bulk mRNAseq analysis results for GT and S samples**

A-B. Scatter (MA) and box plot summary on GT sample (SRR7624365, 0h rep 1) and a corresponding $k = 2$ S sample S1₂ (subplot A) and $k = 2$ and $k = 3$ S samples S1₂ and S1₃ (subplot B).

C-D. Scatter (cross) plots comparing the DE amplitude ($\log_2FC$) calculated on 0h vs 12 h samples for the GT vs $k = 2$ S samples S1₂ (subplot C), and $k = 2$ and $k = 3$ S samples S1₂ and S1₃ (subplot D); the colour gradient is $\log_2$ -proportional to the average abundance across all 8 corresponding samples.

E. Upset plot showing intersections between sets of DE genes (0h vs 12h) in GT samples (GT), 2 sets of $k = 2$ S samples (S1₂ and S2₂) and 2 sets of $k = 3$ S samples (S1₃ and S2₃).

F. Upset plot showing intersections between sets of significant terms (adj p-val ≤ 0.05 , BH correction) predicted using gprofiler2 on the corresponding DE genes (see E).

Comparison of H3K4me3 ChIPseq analysis results for GT and S samples.

G-H. Scatter (MA) and box plots showing $\log_2$ abundance against $\log_2FC$ within G) GT sample (SRR7624381, 0h) and a $k = 2$ S sample S1₂ and H) $k = 2$ and $k = 3$ S samples S1₂ and S1₃ for the same sample (SRR7624381, 0h)

I-J. Scatter (cross) plot showing $\log_2FC$ when comparing I) 0h and 12h for GT samples and $k = 2$ S samples S1₂, and J) 0h and 12h for $k = 2$ and $k = 3$ S samples S1₂ and S1₃, coloured by the average amplitude across all 4 corresponding samples.

K) Violin and box plots showing distribution of ratio of peak lengths between 2 samples ($\frac{\text{sample 1 peak length}}{\text{sample 2 peak length}}$), comparing ground truth (GT), $k = 2$ S sample (S1₂ and S2₂) and $k = 3$ S samples (S1₃).

Effects of across-lane splitting on single-cell data

[sc smartSeq] We exemplify the effect of across-lane sample splitting, on single-cell data, first on a smartSeq2 dataset. The diversity of conditions (donors and time-points, illustrated in Fig 2A) of the Cuomo et al 2020 dataset [30] enables us to incrementally examine the consequences across 6 study cases (Table 1), covering a wide range of experimental situations.

At the sequencing read level, the sample splitting significantly reduces the number of unique reads (Fig 2B, Supp Table 1) i.e. the read diversity decreases due to the expansion (over-representation) of duplicated reads. On an invariant total number of reads, this decrease propagates through to the subsequent alignment and quantification steps without influencing the fraction of mapped reads and the count distribution of samples, respectively (Supplementary Fig 4B,C); nonetheless, this results in a reduction of the number features detected in S samples (Fig 2C). This is a consequence of the loss of low-abundant features, while mid-abundance and high abundance features slightly increase their expression due to the higher read duplication. Significant differences in the number of features (< 0.01 type I error) between simulations and ground truth are detected only for study Case 3; the abundance distributions are not significantly different (in localisation or shape) across the other study cases (Supplementary Tables 1,2). Despite this, the reduction in the number of observed features has an impact on downstream analyses such as partitioning into clusters and cell visualization or identification of cell type markers. For the former, the UMAP topology of cells/ clusters is altered (Supplementary Fig 4A); for the latter, in Fig 2D we observe a generalised reduction in the similarity of marker genes for S samples when compared to marker genes determined on the ground truth sample (overall, 0.55-0.94 of the markers are shared between homologous clusters in GT and S samples). In addition, to assess the variability across simulations, we display summary upset plots of marker genes across clusters and simulations (Fig 2E, Supplementary Fig 5). Although the majority of the DE entries are shared between simulations, unique features are present in most cases; these sum up to 12.3% of total features, on average, across clusters, simulations and study cases.

Overall these results suggest that sample splitting reduces read diversity, which is propagated to downstream analyses, altering the number of features expressed in samples, and potentially altering the biological interpretation of results. Moreover, we highlight the variability across simulations which may introduce an additional degree of irreproducibility across replicates.

Figure 2: Sample splitting across sequencing lanes introduces variability that is propagated to downstream analysis in Smart-seq2 scRNA-seq data. A) UMAP representation of the Cuomo dataset. Cells are coloured according to their donor (left), time-point (central) and cluster (right). For the different study cases considered we show: B) Number of Unique Reads distribution for the ground truth and simulations. C) Number of features distribution for the ground truth and simulations. D) Cluster similarity for each study case as evaluated using the Jaccard Similarity Index calculated on the set of differentially expressed features obtained per cluster. The values across the 3 simulations were averaged using the geometric mean. E) Variability of differentially expressed feature set and number of unique DE features across simulations for cluster 0 (upper), 1 (central) and 2 (lower) for Study Case 1.

[sc 10x Genomics] The second single-cell case study focuses on a 10x dataset [31]; a similar strategy for generating split samples was applied. Following the alignment and protein-coding gene quantification of the GT and S samples, we observe a high overlap of called cells (i.e. barcodes) resulting from the Cellranger cell calling algorithm (19,453 common barcodes, 166 barcodes unique to ground truth, 13 barcodes unique to simulated for P1; 14,944 barcodes in common, 86 barcodes unique to ground truth, 4 barcodes unique to simulated for P2). After filtering the low-quality cells, the vast majority of barcodes were still common to both versions of the analysis (17,095 barcodes in common, 155 barcodes unique to GT, 5 barcodes unique to simulated for P1; 13,040 barcodes in common, 149 barcodes unique to GT, 2 barcodes unique to simulated for P2). Across both donors, we observe a greater diversity of UMIs and genes in GT compared to simulated samples (Supplementary Fig 9).

Following SLM clustering [45] on the nearest-neighbor graph, 19 clusters were found on GT and 18 on the simulation of P1 (Fig 3A,B). Element-centric clustering comparison [46] was used to evaluate the per-cell clustering similarity on the set of barcodes common to GT and simulated samples (Fig 3C), revealing high similarity ($ECS_{\geq 0.6}$) for 91% of cells (the top of the GT UMAP) and especially the for the island on the upper right (cluster 4 in both clusterings); low similarity ($ECS_{\leq 0.2}$) was observed for 9.0% of cells (mainly bottom of the UMAP); no cells had intermediate ECS ($0.2 < ECS < 0.6$). Notably, an island of cells at the lower right of the GT UMAP, corresponding to cluster 15 in GT, disappeared in the split lane simulation (S samples).

To investigate whether the disappearing island was a reproducible effect of the lane splitting simulations, $n = 3$ S samples were generated. The S samples were consistent in terms of UMAP topography and SLM clustering results (Supplementary Fig 7A-F). Across all S samples the "island cells" were scattered in the larger body of cells in each simulated replicate (Supplementary Fig 8E-H). Conversely, in UMAPs generated on GT across 4 random seeds, the island of cells was consistently detached from the larger body of cells (Supplementary Fig 8A-D). Furthermore, the fraction of within-island nearest neighbors for the island cells was quantified in the GT and S samples (Fig 3E); while the island cells had each other as neighbors in GT, this is no longer observed in the simulations, suggesting the loss of some transcriptional heterogeneity during the lane splitting.

The marker genes that distinguish the vanishing island from the rest of the cells are SPP1, SRGN, SOCS2, ALDH1A1, AREG and HIST1H1C (Supplementary Fig 10B); in particular SPP1 appears highly specific to the island. Upon recalculating the PCA, and subsequently the UMAP, on the set of abundant genes with SPP1 excluded, the island is absorbed into the wider body of cells (Supplementary Fig 10C). SPP1 is identified as a highly variable gene in GT by SCTransform (5th most variable gene in P1; 52nd most variable gene in P2). In S samples, marked differences in gene variance can be observed when compared to GT (Supplementary Fig 10E-F), leading to the downstream consequences on dimensionality reductions and clustering results.

To further investigate the consequences of lane splitting clustering variability on cluster markers, typically used to infer the identity of cells, the per-cell Jaccard similarity index (JSI) was calculated (Fig 3D). Certain regions of cells, such as the middle right of the UMAP, displayed high JSI, indicating they would be interpreted similarly in

the GT and S analyses. Other regions exhibited lower JSI; the vanishing island (cluster 15) had lower JSI than the retained island (cluster 4). These results suggest that the former could be interpreted inconsistently, depending on whether the library was split across lanes or not. Similar conclusions (clustering and marker gene variability) can be drawn for P2 across GT and S samples (Supplementary Fig 6); as for P1, consistency across P2 simulation UMAPs was observed (Supplementary Fig 7G-L).

Figure 3: Lane splitting induces variability in clustering of P1 10x scRNA-seq data.

A-B: UMAP representations of GT (A) and S (B) samples with colors indicating SLM clusters calculated on nearest-neighbor graphs; 19 clusters are found in GT and 18 in the S samples.

C. Element-centric clustering similarity, highlighted using the colour gradient, reveals differences between clusterings at the bottom of the UMAP, especially for the lower right island of cells (vanishing island). D. Jaccard similarity of cluster markers across GT and S clusterings suggests differences in cell types inferred from the data. E. Fraction of within-island nearest neighbors for cells in the vanishing island in GT and 3 S samples. The island cells are no longer discrete from the larger body of cells in S samples.

Discussion

Effects of splitting on read diversity and levels of noise

The main consequence of across-lane sample splitting is the variation in read diversity; this propagates onto transcript coverage and quantification. To assess the variation, we focused on the transcript complexity (defined as the ratio of unique to total reads [48], calculated per transcript), exemplified on the SRR7624365 sample (bulk mRNAseq, 0h rep 1). We compared the complexities, per transcript, of the GT and S samples (Fig 4A); each point corresponds to a gene, the colour gradient is proportional to the $\log_2(\text{abn})$. We observe a consistent trend for medium-high abundant genes, across a wide complexity range $[0, 0.75]$, highlighting the higher complexity (i.e. more diverse reads) of the GT sample. Also noticeable is the high variability in complexity for the low abundance genes, presented on the MA plots in the bulk section of the Results, and the localisation of the low abundance cloud under the equal complexity diagonal, enforcing the previous conclusion across all abundances.

Yet another side effect of the variation in reads diversity is the quantification of noise across the samples [14]; we focused on the transcript-based approach in noisyR since it is directly influenced by the robustness of transcript coverage. In Fig 4B, we illustrate the variation of point-to-point PCC vs the variation in abundance for the GT and 2 iterations of $k = 2$ (S1₂, S2₂) and $k = 3$ (S1₃, S2₃). The wide and low PCC distributions correspond to higher levels of noise; the distributions become higher and tighter for medium to high abundance genes. For medium abundance genes, the GT distributions are systematically higher than $k = 2$ and $k = 3$, suggesting that the splitting increases the level of noise and thus interferes with the detection of DE genes or other downstream analyses, as illustrated in the Results section. The level of noise in the simulated samples is also variable and generally higher in $k = 3$ samples than $k = 2$, although the distribution of PCC is wider and lower for all S samples than GT.

Effects of varying the number and proportions of splits

In the Results section, we exemplified the concatenation of $k = 2$ and $k = 3$ equal subsamples. However, in a real-world scenario, we rarely observe an equal number of reads across split samples; in addition, split-designs are occasionally mixed. To further understand the knock-on effects on downstream analyses, we illustrate the effect of variable number and proportions of reads allocated to splits. We used the bulk mRNAseq sample SRR7624365 (0h rep 1) as the starting point and, based on the consistency in conclusions across different inputs, we expect a convergence of results for other bulk or single-cell data. We subsampled the fastq files to 50%, 33%, 25%, 20% and 10% of the total sequencing depth and created S samples by concatenating the 2, 3, 4, 5 and 10 subsamples, respectively. For each case study, the distributions were generated on 10 iterations. We also assessed the sequencing depth co-variate; the results on the full sequencing depth (75.4M) and for samples subsampled to 50M, 25M and 10M reads respectively (without replacement) are presented. The ratios of recovered unique reads (i.e. number of unique reads in the S sample divided by the number of unique reads for the GT sample) across the simulated case

Figure 4: **Summary of split-effects on the number of unique reads and noise.** A) Scatter plot illustrating the complexity ratio (non-redundant to redundant counts) for the GT (x-axis) vs an S (y-axis) sample for SRR7624365 (bulk mRNAseq, 0h rep 1). Each point represents a gene and the colour gradient is proportional to the $\log_2$ abundance. B) Box plot of the PCC binned by abundance for the transcript-based noise removal (noisyR applied to BAM files) corresponding to GT, $k = 2$ and $k = 3$ S samples for 0h rep 1. C) Boxplot showing distributions of ratios of recovered unique reads in the S samples with respect to the k hyper-parameter ($k = 2, 3, 4, 5, 10$); the effect of GT sample sequencing depth is also assessed. The distributions are built on 10 iterations. D) Boxplot showing distribution of ratios of recovered unique reads in the S samples when the proportions of concatenated subsamples are varied (from 55-45 to 90-10). The distributions are built on 10 iterations.

studies are presented in Fig 4C. We observe a decrease in recovered ratios proportional to the number of splits and initial sequencing depth. The decrease in ratios is uniform across sequencing depths, but the recovered ratios also have lower y-intercepts (lower absolute ratios), highlighting that the effect of across lane-splitting is more extreme at low sequencing depths i.e. while the decrease-rate is lower for higher k -values, as number of splits (k) increases, we see a consistent reduction in the number of unique reads. This underlines that splitting across lanes has an adverse effect on the diversity of reads.

Additionally, we assessed the effect of uneven splits across lanes; focusing on $k = 2$ S samples, we varied the split proportions from 55-45 to 90-10. The resulting recovery ratios are shown in Fig 4D. The minimum for the recovery ratios is achieved for the 50-50 proportions; the recovery ratios gradually increases as the larger subsample approaches 100%. The observed increase is small for first few increments and increases more rapidly as proportions approach 100-0. This illustrates, from yet another angle, the variation in reads diversity.

Key points

On various sequencing datasets, bulk and single cell, we observed that the splitting of samples across lanes reduces the diversity of reads, which in turn triggers side effects on quantification (e.g. gene expression for mRNAseq, peak expression for ChIPseq) and auxiliary properties (such as the length of ChIP peaks). The splitting in itself introduces an additional level of variability in terms of robustness and reproducibility; it may pose added difficulties stemming from variable observed number of reads derived from the stochasticity of the sequencing itself. The potential batch effects derived from loading full samples on sequencing lanes can be mitigated through randomisation. We acknowledge that technical circumstances may make splitting unavoidable; our recommendation is a consistency in sequencing setup across all samples in an experiment.

290 Acknowledgments

291 All authors acknowledge the constructive feedback from the Core Bioinformatics group and the support by the Core
 292 grant awarded to the Wellcome-MRC Cambridge Stem Cells Institute. This research was funded in whole, or in part,
 293 by the Wellcome Trust [203151/Z/16/Z] and the UKRI Medical Research Council [MC_PC_17230]. We also thank
 294 Dr Chris Ellis for discussion on the splitting applied on 10x Genomics data. IM acknowledges early discussions of
 295 the idea that led to is project with Dr Rachel Rusholme-Pilcher (Earlham Institute, NRP, UK).

296 Author contributions

297 EW performed the bulk analyses (mRNAseq and ChIPseq); RCG performed the smartSeq analyses; AS performed
 298 the 10x Genomics analyses. IM proposed and supervised the project. All authors wrote the manuscript, and read
 299 the final draft.

References

- [1] Stark R, Grzelak M, and Hadfield J. RNA sequencing: the teenage years. *Nature Reviews Genetics*, 20, 07 2019.
- [2] Yandell M and Ence D. A beginner’s guide to eukaryotic genome annotation. *Nature Reviews Genetics*, 13:329–42, 04 2012.
- [3] Steward C, Parker A, Minassian B, et al. Genome annotation for clinical genomic diagnostics: Strengths and weaknesses. *Genome Medicine*, 9, 05 2017.
- [4] Salzberg S. Next-generation genome annotation: We still struggle to get it right. *Genome Biology*, 20, 12 2019.
- [5] Conesa A, Madrigal P, Tarazona S, et al. A survey of best practices for RNA-seq data analysis. *Genome Biology*, 17, 01 2016.
- [6] Oshlack A, Robinson M, and Young M. From RNA-seq reads to differential expression results. *Genome Biology*, 11:220, 12 2010.
- [7] Lightbody G, Haberland V, Browne F, et al. Review of applications of high-throughput sequencing in personalized medicine: barriers and facilitators of future progress in research and clinical application. *Briefings in Bioinformatics*, 20(5):1795–1811, 06 2019.
- [8] Lücken M and Theis F. Current best practices in single-cell RNA-seq analysis: a tutorial. *Molecular Systems Biology*, 15:e8746, 06 2019.
- [9] McGuire A, Gabriel S, Tishkoff S, et al. The road ahead in genetics and genomics. *Nature Reviews Genetics*, 21:1–16, 08 2020.
- [10] Stupnikov A, Tripathi S, de Matos Simoes R, et al. samExploreR: exploring reproducibility and robustness of RNA-seq results based on SAM files. *Bioinformatics*, 32(21):3345–3347, 07 2016.
- [11] Schurch N, Schofield P, Gierliński M, et al. How many biological replicates are needed in an RNA-seq experiment and which differential expression tool should you use? *RNA*, 22, 03 2016.
- [12] Oberg A, Bot B, Grill D, et al. Technical and biological variance structure in mrna-seq data: life in the real world. *BMC Genomics*, 13:304, 07 2012.
- [13] Kim BS, Lee E, and Kim J. *Analysis of technical and biological variability in single-cell RNA sequencing*, volume 1935, pages 25–43. Humana Press, 01 2019.
- [14] Moutsopoulos I, Maischak L, Lauzikaite E, et al. noisyR: Enhancing biological signal in sequencing datasets by characterising random technical noise. *bioRxiv*, 2021.

- 328 [15] Ma X, Shao Y, Tian L, et al. Analysis of error profiles in deep next-generation sequencing data. *Genome*
329 *Biology*, 20, 03 2019.
- 330 [16] Ross M, Russ C, Costello M, et al. Characterizing and measuring bias in sequence data. *Genome Biology*,
331 14:R51, 05 2013.
- 332 [17] Hicks S, Townes FW, Teng M, and Irizarry R. Missing data and technical variability in single-cell RNA-
333 sequencing experiments. *Biostatistics (Oxford, England)*, 19, 11 2017.
- 334 [18] Reuter J, Spacek D, and Snyder M. High-throughput sequencing technologies. *Molecular cell*, 58:586–597, 05
335 2015.
- 336 [19] Eraslan G, Simon L, Mircea M, et al. Single-cell RNA-seq denoising using a deep count autoencoder. *Nature*
337 *Communications*, 10:390, 01 2019.
- 338 [20] Chazarra-Gil R, Dongen S, Kiselev V, et al. Flexible comparison of batch correction methods for single-cell
339 RNA-seq using BatchBench. *Nucleic Acids Research*, 49, 02 2021.
- 340 [21] Srivastava A, Malik L, Sarkar H, et al. Alignment and mapping methodology influence transcript abundance
341 estimation. *Genome Biology*, 21:239, 09 2020.
- 342 [22] Dillies M, Rau A, Aubert J, et al. A comprehensive evaluation of normalization methods for Illumina high-
343 throughput RNA sequencing data analysis. *Briefings in Bioinformatics*, 14, 09 2012.
- 344 [23] Mccarthy D, Chen Y, and Smyth G. Differential expression analysis of multifactor SRNA-Seq experiments with
345 respect to biological variation. *Nucleic Acids Research*, 40:4288–97, 02 2012.
- 346 [24] Love M, Huber W, and Anders S. Moderated estimation of fold change and dispersion for RNA-Seq data with
347 DESeq2. *Genome Biology*, 15:550, 12 2014.
- 348 [25] Svensson V, Natarajan K, L Ly, et al. Power analysis of single cell RNA-sequencing experiments. *Nature*
349 *Methods*, 14, 03 2017.
- 350 [26] Nakato R and Shirahige K. Recent advances in ChIP-seq analysis: From quality management to whole-genome
351 annotation. *Briefings in Bioinformatics*, 18:bbw023, 03 2016.
- 352 [27] Chung D, Kuan PF, Li B, et al. Discovering Transcription Factor Binding Sites in Highly Repetitive Regions
353 of Genomes with Multi-Read Analysis of ChIP-Seq Data. *PLoS Computational Biology*, 7:e1002111, 07 2011.
- 354 [28] Dal Molin A and Camillo B. How to design a single-cell RNA-sequencing experiment: Pitfalls, challenges and
355 perspectives. *Briefings in Bioinformatics*, 20, 01 2018.
- 356 [29] Yang P, Humphrey SJ, Cinghu S, et al. Multi-omic profiling reveals dynamics of the phased progression of
357 pluripotency. *Cell Systems*, 8(5):427–445.e10, 2019.

- 358 [30] Cuomo A, Seaton D, McCarthy D, et al. Single-cell RNA-sequencing of differentiating iPS cells reveals dynamic
359 genetic effects on gene expression. *Nature Communications*, 11, 02 2020.
- 360 [31] Mende N, Bastos HP, Santoro A, et al. Quantitative and molecular differences distinguish adult human medullary
361 and extramedullary haematopoietic stem and progenitor cell landscapes. *bioRxiv*, 2020.
- 362 [32] Yates AD, Achuthan P, Akanni W, et al. Ensembl 2020. *Nucleic Acids Research*, 48(D1):D682–D688, 11 2019.
- 363 [33] Dobin A, Davis CA, Schlesinger F, et al. STAR: ultrafast universal RNA-seq aligner. *Bioinformatics*, 29(1):15–
364 21, 10 2012.
- 365 [34] Liao Y, Smyth GK, and Shi W. featureCounts: an efficient general purpose program for assigning sequence
366 reads to genomic features. *Bioinformatics*, 30(7):923–930, 11 2013.
- 367 [35] Bolstad BM, Irizarry RA, Åstrand M, et al. A comparison of normalization methods for high density oligonu-
368 cleotide array data based on variance and bias. *Bioinformatics*, 19:185–193, 01 2003.
- 369 [36] Robinson M, Mccarthy D, and Smyth G. edgeR: a Bioconductor package for differential expression analysis of
370 digital gene expression data. *Bioinformatics*, 26:139–140, 01 2010.
- 371 [37] Raudvere U, Kolberg L, Kuzmin I, et al. g:Profiler: a web server for functional enrichment analysis and
372 conversions of gene lists (2019 update). *Nucleic Acids Research*, 47(W1):W191–W198, 05 2019.
- 373 [38] Zhang Y, Liu T, Meyer C, et al. Model-based analysis of ChIP-seq (MACS). *Genome Biology*, 9:R137, 10 2008.
- 374 [39] Andrews S, Krueger F, Segonds-Pichon A, et al. FastQC. Babraham Institute, January 2012.
- 375 [40] Ewels P, Magnusson M, Lundin S, et al. MultiQC: summarize analysis results for multiple tools and samples in
376 a single report. *Bioinformatics*, 32(19):3047–3048, 06 2016.
- 377 [41] Stuart T, Butler A, Hoffman P, et al. Comprehensive integration of single cell data. *Cell*, 11 2018.
- 378 [42] Hafemeister C and Satija R. Normalization and variance stabilization of single-cell RNA-seq data using regu-
379 larized negative binomial regression. *Genome Biology*, 20, 12 2019.
- 380 [43] Beckers M, Mohorianu I, Stocks M, et al. Comprehensive processing of high throughput small RNA sequenc-
381 ing data including quality checking, normalization and differential expression analysis using the UEA sRNA
382 Workbench. *RNA*, 23:rna.059360.116, 03 2017.
- 383 [44] Zheng G, Terry J, Belgrader P, et al. Massively parallel digital transcriptional profiling of single cells. *Nature*
384 *Communications*, 8:14049, 01 2017.
- 385 [45] Waltman L and van Eck NJ. A smart local moving algorithm for large-scale modularity-based community
386 detection. *European Physical Journal B*, 86, 08 2013.

- 387 [46] Gates A, Wood I, Hetrick W, et al. Element-centric clustering comparison unifies overlaps and hierarchy.
388 *Scientific Reports*, 9:8574, 06 2019.
- 389 [47] Shahsavari A and Mohorianu I. *ClustAssess: Tools for Assessing Clustering*, 2021. R package version 0.1.1.
- 390 [48] Mohorianu I, Schwach F, Jing R, et al. Profiling of short RNAs during fleshy fruit development reveals stage-
391 specific sRNAome expression patterns. *The Plant journal : for cell and molecular biology*, 67:232–46, 03 2011.