

# Depth normalization of small RNA sequencing: using data and biology to select a suitable method

Yannick Düren<sup>1,2</sup>, Johannes Lederer<sup>1</sup> and Li-Xuan Qin<sup>1,2,\*</sup>

<sup>1</sup>Department of Mathematical Statistics, Ruhr-University Bochum, Universitätsstraße 150, 44801 Bochum, Germany and <sup>2</sup>Department of Epidemiology and Biostatistics, Memorial Sloan Kettering Cancer Center, New York, NY 10065, USA

Received June 02, 2021; Revised January 03, 2022; Editorial Decision January 18, 2022; Accepted February 08, 2022

## ABSTRACT

Deep sequencing has become one of the most popular tools for transcriptome profiling in biomedical studies. While an abundance of computational methods exists for ‘normalizing’ sequencing data to remove unwanted between-sample variations due to experimental handling, there is no consensus on which normalization is the most suitable for a given data set. To address this problem, we developed ‘DANA’—an approach for assessing the performance of normalization methods for microRNA sequencing data based on biology-motivated and data-driven metrics. Our approach takes advantage of well-known biological features of microRNAs for their expression pattern and chromosomal clustering to simultaneously assess (i) how effectively normalization removes handling artifacts and (ii) how aptly normalization preserves biological signals. With DANA, we confirm that the performance of eight commonly used normalization methods vary widely across different data sets and provide guidance for selecting a suitable method for the data at hand. Hence, it should be adopted as a routine preprocessing step (preceding normalization) for microRNA sequencing data analysis. DANA is implemented in R and publicly available at <https://github.com/LXQin/DANA>.

## INTRODUCTION

Deep sequencing is prone to systematic non-biological artifacts that arise from variations in experimental handling, similar to other genomics technologies such as microarrays (1,2). Consequently, a critical first step in the analysis of transcriptome sequencing data is to ‘normalize’ the data so that data from different sequencing runs are comparable (3–5). One major source of such handling effects comes from the depth of coverage — defined as the average number of reads per molecule (6).

A plethora of analytic methods for depth normalization have been proposed, including methods based on re-scaling (7–10) and methods based on regression (3,11,12). Different normalization methods may lead to different analysis results, and no method has been found to work systematically best in studies comparing the performance of these methods (10,13–16). Rather, their performance strongly depends on the data under study (2,17). Currently, a method is often chosen by the analyst based on personal preference and convenience rather than objective criteria governed by the data.

In this study, we introduce a data-driven and biology-motivated approach to objectively guide the selection of depth normalization methods for the data at hand. We call our novel approach ‘DATA-driven Normalization Assessment’ (DANA). DANA’s goal is to identify a method that maximally removes depth variations due to disparate experimental handling (‘handling effects’) while minimally impacting true biological signals (‘biological effects’). Its overall three-step strategy is to first define two sets of control markers for capturing each of the two types of variations, then compute statistical measures for quantifying these variations and their change before *versus* after normalization, and, lastly, use numeric metrics and graphical tools for summarizing these measures across each set of control markers.

We apply this concept to the sequencing of microRNAs (miRNAs), a prevalent class of small RNAs that play an important regulatory role of gene expression in the cell (18,19). Two key biological features of miRNAs are that (i) they tend to be expressed in an on-off manner where only a subset of miRNAs are expressed in a given sample and (ii) a subset of them can be organized into polycistronic clusters that tend to be co-regulated and hence co-expressed (20–23). To exploit these known features, we first define a set of negative control markers as those miRNAs that are poorly expressed (i.e., markers with low mean expression) reflecting primarily handling effects. When handling effects exist in the data, they manifest as high positive correlations among these negative controls, simply due to the shared handling effects. We further define a set of positive control markers

\*To whom correspondence should be addressed. Tel: +1 646 888 8251; Fax: +1 646 888 0010; Email: [qinl@mskcc.org](mailto:qinl@mskcc.org)

as the collection of miRNAs that are well-expressed (i.e., markers with high mean expression) and belong to polycistronic clusters. The shared biological effects lead to high positive correlations among these positive controls regardless of handling effects. Our DANA approach does not require any additional study design (such as balanced sample-to-sequencing-batch assignment) or reference data (such as reliable spike-in markers) that is not readily available. To the best of our knowledge, DANA is the first approach to achieve a purely data-driven assessment and selection of depth normalization method for miRNA sequencing data. While this concept can be generalized and applied to other molecules, such as mRNAs, the definition of control markers may not be as straightforward as for miRNAs and will be the subject of future work.

We benchmark the performance of DANA using a unique pair of miRNA sequencing data sets for the same set of tumor samples, which were previously collected at Memorial Sloan Kettering Cancer Center (16). The first data set was sequenced using uniform handling to minimize handling effects and balanced sample-to-library-assignment to cancel out any residual handling effects in group comparison. For the same set of samples, a second data set was collected over the years in the order of tumor sample collection and exhibited excessive depth variations. We use this pair of data sets to validate our definition of control markers, choice of correlation measures and summary metrics, and the performance of the overall DANA approach. We further validate and demonstrate the use of our approach by applying it to two data sets from The Cancer Genome Atlas (TCGA).

## MATERIALS AND METHODS

We first present our three-step DANA approach, next introduce the paired data sets from MSK for its benchmarking, and then describe two data sets from TCGA for further validation and demonstration. The DANA method is implemented in R and open source (<https://github.com/LXQin/DANA>).

### DANA: A data-driven and biology-motivated normalization selection method

#### Step I. Definition of negative and positive controls.

**Concept.** Following the ideas in (24), we define poorly-expressed markers as those with mean abundance (that is, read count) in an interval  $[\ell^-, u^-]$  and well-expressed markers as those with mean abundance over a cutoff  $\ell^+$ . It is well known that miRNAs tend to be expressed in an on-off manner where only a subset of miRNAs are expressed in a given sample (21,25). Hence, poorly-expressed markers mainly reflect those markers that were expressed with read counts driven by handling effects. These poorly-expressed markers serve as negative controls. It is further known that miRNAs belonging to a shared polycistronic cluster tend to be co-regulated and, hence, co-expressed (20–23). Due to this co-expression, we can consequently expect a high, biology-driven correlation between markers located in a mutual polycistronic cluster, regardless of handling effects. Therefore, we use well-expressed markers that belong to

polycistronic clusters with at least two well-expressed members as positive controls reflecting shared biological effects. For simplicity and reproducibility, we define miRNAs to be in a mutual polycistronic cluster if their hairpins are separated by  $<10$  kb on the chromosome. This corresponds to the cluster definition on miRBase (22).

**Cutoff selection.** The cutoffs  $\ell^-, u^-$  and  $\ell^+$  should be chosen based on the empirical distribution of the data. We set  $\ell^-$  such that selected poorly-expressed miRNAs show at least mild expression based on the mean read count histogram. For all data considered in this paper, we observed that  $\ell^- = 2$  yielded reasonable results and that different cutoff choices only mildly affect the subsequent analysis. We choose  $u^-$  and  $\ell^+$  such that each control group consists of a sufficient number of markers, where we recommend 75–150 markers per group for balancing statistical stability and computational convenience. Note that if the data contains many subtypes or the number of samples is highly unbalanced across subtypes, it might be necessary to take special care during the selection of negative controls. For example, if a marker is expressed only in a single subtype but non-expressed in all others, it may be erroneously classified as a negative control marker by simply taking the mean expression. In such cases, we instead recommend selecting negative control markers based on the maximum across subtype means.

Henceforth, we denote the number of positive and negative controls by  $p^+$  and  $p^-$ , respectively. For ease of notation, we re-index the positive controls by  $\mathcal{C}^+ := \{1, \dots, p^+\}$  and the negative controls by  $\mathcal{C}^- := \{1, \dots, p^-\}$ . Furthermore, we denote the set of all pairs of distinct positive control miRNAs  $(i, j)$  that are in a mutual polycistronic cluster by

$$\mathcal{C}^+ := \{(i, j) \in (\mathcal{C}^+)^2 \mid (i < j) \text{ and} \\ (i \text{ and } j \text{ are in a mutual polycistronic cluster})\}.$$

#### Step II. Statistical measures of between-marker correlation.

**Concept.** To measure the level of correlation for a pair of control markers, we estimate either their marginal or partial correlations, depending on the control type. For negative controls, we quantify the overall strength of inter-marker correlations, which are primarily due to shared handling effects, using Pearson marginal correlations. For positive controls, on the other hand, we capture only the direct co-expression relation for each pair of markers, upon removing any spurious correlation due to co-expression with other positive controls, using partial correlations. Note that for the computation of empirical correlations, the sample size  $n$  must not be overly small. In our experiments, we have observed that  $n \geq 20$  is sufficient for typical miRNA sequencing data.

**Statistical definition of partial correlation.** For estimating partial correlations in positive controls, we assume a parametric distribution for the  $\log_2$ -transformed read counts, where by ‘ $\log_2$ -transformation’, we generally refer to the  $\log_2(\cdot + 1)$  function so that zero counts are assigned a value of zero after the transformation. More specifically, we use a multivariate normal distribution as it is widely used for

modeling (log2-transformed) sequencing data (26), and as it is a standard distribution in statistics with many tools readily available. Our pipeline also allows for other distributions, such as Poisson or negative binomial for the count data, but the estimation would be much more challenging (27). Under the assumption of normal-distributed data, the Hammersley–Clifford theorem (see for instance (28)) establishes a direct link of partial correlations of  $p$  variables to the entries of their precision matrix  $\Theta = \Sigma^{-1} \in \mathbb{R}^{p \times p}$ , the inverse of the covariance matrix  $\Sigma \in \mathbb{R}^{p \times p}$  of the data. For any two variables indexed by  $i$  and  $j$ , respectively, their partial correlation  $\rho_{i,j}$  follows the relation

$$\rho_{i,j} = -\frac{\theta_{i,j}}{\sqrt{\theta_{i,i}\theta_{j,j}}},$$

where  $\theta_{i,j}$  is the  $(i, j)$ th entry of the precision matrix  $\Theta = (\theta_{i,j})_{1 \leq i,j \leq p}$ . Thus, the partial correlation structure can be estimated using well-established statistical methods for precision matrix estimation such as neighborhood selection (29), graphical lasso (30), or FastGGM (31). In this study, we use the popular neighborhood selection method proposed by Meinshausen and Bühlmann (29) and calibrate its tuning parameter using Bayesian Information Criterion. We, furthermore, compare the results of the chosen method to all aforementioned methods for precision estimation.

**Estimation of correlations before and after normalization.** We estimate partial correlations in positive controls and Pearson correlations in negative controls using un-normalized log-counts as well as normalized log-counts for each normalization method under study. We denote the estimated partial correlations in positive controls for un-normalized log-counts by  $\rho^+ := (\rho_{i,j}^+)_{1 \leq i,j \leq p^+}$ , and for an arbitrary normalization method by  $\rho^{\text{norm},+} := (\rho_{i,j}^{\text{norm},+})_{1 \leq i,j \leq p^+}$ . In parallel, we denote Pearson correlations in negative controls by  $\rho^- := (\rho_{i,j}^-)_{1 \leq i,j \leq p^-}$ , and  $\rho^{\text{norm},-} := (\rho_{i,j}^{\text{norm},-})_{1 \leq i,j \leq p^-}$ , respectively.

**Step III. Numerical metrics for summarizing correlation across markers.** The performance of a normalization method is evaluated through a comparison of the estimated correlations before *versus* after normalization in positive controls ( $\rho^+$  with  $\rho^{\text{norm},+}$ ) and in negative controls ( $\rho^-$  with  $\rho^{\text{norm},-}$ ). We introduce two numeric metrics, one for each control type.

**A numeric metric for the reduction of handling effects.** An effective normalization method should maximally remove marginal correlations among negative controls, which are most likely caused by shared handling effects. Ideally, correlations in the normalized data should be centered around zero with low variance. Therefore, we propose to compute the mean-squared discrepancy of correlations from zero

$$v_0 := \frac{1}{p^-(p^- - 1)} \sum_{i=1}^{p^-} \sum_{\substack{j=1 \\ j \neq i}}^{p^-} (\rho_{i,j}^- - 0)^2,$$

for un-normalized and, analogously,  $v_0^{\text{norm}}$  for normalized data. We use the metric

$$\text{mscr}^- := \frac{v_0 - v_0^{\text{norm}}}{v_0}, \quad (1)$$

i.e., the relative difference of the variance of correlations among negative controls before and after normalization to quantify the reduction of handling effects through normalization. We call this metric the ‘mean-squared correlation reduction in negative controls,’ where a high  $\text{mscr}^-$  indicates an effective removal of handling effects in the data.

**A numeric metric for the preservation of biological effects.** For each pair of positive control miRNAs that are located in a mutual polycistronic cluster, an ideal normalization method should minimally affect their co-expression. Mathematically speaking, this means that the relationship between  $\rho_{i,j}^+$  and  $\rho_{i,j}^{\text{norm},+}$  for all same-cluster positive controls  $(i, j) \in \mathcal{C}^+$  should be linear with slope 1. As such, we measure this agreement by the concordance correlation coefficient (32). We denote the vector of correlations in  $\mathcal{C}^+$  by  $(\rho^+)_{\mathcal{C}^+}$  for un-normalized data, and by  $(\rho^{\text{norm},+})_{\mathcal{C}^+}$  for normalized data. We further denote the concordance correlations coefficient of two covariates  $x_1, x_2$  by  $\text{CCC}[x_1, x_2]$ . We use the metric

$$\text{cc}^+ := \text{CCC}[(\rho^+)_{\mathcal{C}^+}, (\rho^{\text{norm},+})_{\mathcal{C}^+}], \quad (2)$$

that is, the ‘concordance correlation coefficient of the within-cluster partial correlations among positive controls before and after normalization,’ to quantify the preservation of biological signals. A close-to-1  $\text{cc}^+$  indicates an apt preservation of biological signals in the data.

**Guidelines for method selection.** An optimal normalization maximally removes handling effects (high  $\text{mscr}^-$ ) while keeping biological signals intact ( $\text{cc}^+$  close to 1). However, in most cases, there is no clear ‘best’ method with maximal  $\text{mscr}^-$  and maximal  $\text{cc}^+$ . Therefore, one should aim for the best possible trade-off between the proposed statistics for negative and positive controls with an emphasis on keeping biological signals intact ( $\text{cc}^+$  close to 1). The two metrics can be easily assessed by plotting the metrics in a scatter plot for each normalization method under study, where a preferable method should be located towards the top-right quadrant of the plot.

### Graphical tools

We introduce several graphical tools for examining the definition of positive and negative controls, providing empirical evidence that supports the choice of DANA’s correlation measures and summary metrics, and comparing the two metrics among the normalization methods for their performance assessment. First, a mean count histogram and mean-standard deviation plot should be used to visually examine the un-normalized data and define positive and negative control markers based on the observed count *versus* variation distribution. Second, we use marginal correlation histograms, partial correlation heatmaps, and partial correlation scatter plots for (i) visually examining the estimated

correlations and (ii) further justifying the choice of the summary metrics  $cc^+$  and  $mscr^-$ . Lastly, we use scatter plots of DANA's two metrics for comparing the normalization methods under assessment for guiding the method selection.

**Mean-standard deviation plot.** Mean-standard deviation plots are scatter plots of marker-specific standard deviations *versus* marker-specific means after  $\log_2$ -transformation. For miRNA sequencing data, we typically observe close-to-zero standard deviation for non-expressed markers. For increasing marker-specific mean, the standard deviation typically increases and 'fans out.' We recommend that for negative controls, the lower bound  $\ell^-$  is tuned such that selected miRNAs reflect a non-zero and heterogeneous distribution of standard deviations capturing miRNAs that reflect handling effects.

**Mean count histogram.** After a  $\log_2$ -transformation of the read count data, we compute the mean read count for each miRNA across all samples. The distribution and range of reads in the data are then visualized using a histogram plot for the mean read abundance. This way, miRNAs that are essentially non-expressed can be excluded, and poorly- or well-expressed miRNAs can be easily identified.

**Marginal correlation histogram.** We compare the frequency of Pearson correlations among negative controls before *versus* after normalization using a histogram. The distribution of  $\rho^-$  and  $\rho^{\text{norm},-}$  is shown by binning correlation strengths and counting the number of observations in each bin. Moreover, we use the correlation histogram for comparing correlation strengths between different data sets, both before and after normalization. An effective normalization method should center the correlations around zero with small variance.

**Partial correlation heatmap.** The partial correlation heatmap shows the matrix of pairwise partial correlations for positive controls  $\mathcal{C}^+$  alongside the polycistronic clustering. Its upper triangular part shows the strength of the partial correlation, while the lower part indicates whether the pair of miRNAs is located in a mutual polycistronic cluster. Using partial correlation heatmaps, we visually identify within-cluster correlations in the data and compare them before and after normalization. The heatmap also allows the comparison of within-cluster correlations with between-cluster correlations, both before and after normalization.

**Partial correlation scatter plot.** The concordance of partial correlations in positive controls before normalization  $\rho^+$  *versus* after normalization  $\rho^{\text{norm},+}$  is shown using a scatter plot. For each pair of distinct, clustered, positive control miRNAs  $(i, j) \in \mathcal{C}^+$ , their partial correlation  $(\rho_{i,j}^+, \rho_{i,j}^{\text{norm},+})$  is plotted in a scatter plot. An ideal normalization method should minimally affect  $\rho_{i,j}^+$ , and hence, each  $(\rho_{i,j}^+, \rho_{i,j}^{\text{norm},+})$  should ideally lay on a diagonal with slope 1. Combined for all  $(i, j) \in \mathcal{C}^+$ , this concordance is quantified by our metric  $cc^+$ .

## MSK sarcoma data

**A pair of data sets for soft tissue sarcoma.** We have previously collected two data sets for the *same set* of tumor samples at MSK. The tumor samples were 27 myxofibrosarcoma (MXF) samples and 27 pleomorphic malignant fibrous histiocytoma (PMFH) samples, which were all from newly diagnosed, previously untreated, primary tumors collected at MSK between 2000 and 2012. The first data set was collected using uniform handling to minimize data artifacts and balanced sample-to-library-assignment (via the use of blocking and randomization) to balance any residual artifacts with the tumor groups under comparison. For extra quality assurance, we added two pooled samples shared across all libraries and 10 calibrators spiked-in at fixed concentrations. For the same set of samples, a second data set was collected without making use of such a careful study design, resulting in unwanted depth variations. The number of observed miRNAs is 1033 and a detailed description of the data was previously reported in (16). Throughout this paper, we refer to the uniformly-handled data set as the 'benchmark' data set and the non-uniformly-handled data set as the 'test' data set, following the notation used in (16). The collection of both data sets is called the 'MSK data sets.' The benchmark data set serves a two-fold purpose: (i) it provides a baseline reference for assessing the choice of the control markers and their correlation measures, and (ii) it offers a gold standard of differential expression identification (comparing MXF and PMFH) for checking whether the chosen normalization method, which our approach deems optimal for the test data, leads to the closest differential expression assessment.

**Normalization methods assessed in the test data set.** Our data-driven normalization selection method generally works for *any* between-sample normalization method that generates non-negative normalized data. As proof of concept, we examined eight commonly used normalization methods that were studied in (16), including six scaling-based methods and two regression-based methods. A scaling-based method calculates a scaling factor based on the data for each sample and divides its counts by this factor. A regression-based method can be non-parametric or parametric: non-parametric methods are based on, for example, a quantile-quantile plot; parametric methods are based on a linear regression, which typically adopts a covariate representing depth variation (by using a known batch variable or deriving a surrogate batch variable from the data) and including this covariate in a regression framework for the analysis of differential expression. In this paper, we examined the following popular methods: Total Count (TC) (10), Upper Quartile (UQ) (13), Median (Med) (10), Trimmed Mean of M-values (TMM) (8), DEseq (9), PoissonSeq (33), Quantile Normalization (QN) (34), and Remove Unwanted Variation (RUV) (3) with its three variants RUVg, RUVs, and RUVr. For a description of these methods, we refer readers to (16). The DANA approach is applied to the test data set to compare the performance of these normalization methods and guide the choice of a most suitable one.

**Differential expression assessment in both data sets.** Similar to (16), we first assessed evidence of differential expression in the benchmark data (without normalization) using the voom method (26,35) with a  $P$ -value cutoff of 0.01, serving as a ‘gold standard.’ We then assessed differential expression in the test data before and after normalization and compared it with the gold standard. This comparison was summarized numerically using the sensitivity (true positive rate) and the positive predictive value ( $1 - \text{false discovery rate}$ ), at the risk of abusing the terminology.

### TCGA-UCEC data

We further support and validate our approach with two data sets that are partially paired from the TCGA Uterine Corpus Endometrial Carcinoma (TCGA-UCEC) project (36). We constructed the two data sets so that they each contain 48 samples, 22 of which were of endometrioid subtype (END) and 26 of serous subtype (SER), and so that both data sets have 24 samples in common. The 48 samples of the first data set were processed in a single batch (batch number 228.63.0), which we will refer to as the ‘single-batch’ data. The second data set is composed of 24 samples (11 END and 13 SER) from the first data set and 24 samples (11 END and 13 SER) from other batches, which we will refer to as the ‘mixed-batch’ data. Sequencing counts are available for 1848 miRNAs in each sample. A detailed description of these data sets is deferred to Section 2.1 of the Supplementary Material.

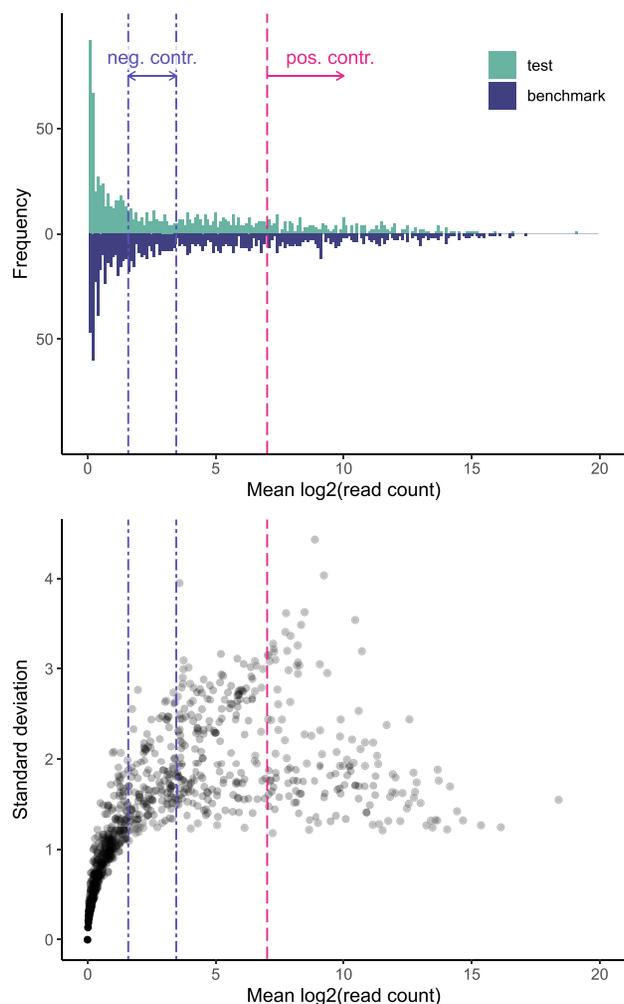
It is reasonable to assume that the mixed-batch data set contains significantly higher variation due to handling effects compared to the single-batch data set. At the same time, due to the sample overlap and identical sample size, biological differences between the two subtypes should be similar between the two data sets. In other words, the single-batch data can provide a ‘silver standard’ for the differential expression status, and the mixed-batch data can serve as a test data set for normalization methods and their selection. We follow the same steps for normalizing the mixed-batch data and performing differential expression analysis in each data set as for the MSK data.

### Combined TCGA-BRCA and TCGA-UCS data

We lastly demonstrate the necessity and effectiveness of our DANA approach in a TCGA data set *without an available benchmark*. We combined all 166 stage III/IV samples from the TCGA Breast Invasive Carcinoma (TCGA-BRCA) project (37) with all 57 samples from the Uterine Carcinosarcoma (TCGA-UCS) project (38). The total sample size ( $n = 223$ ) is an order of magnitude greater than that for the MSK and TCGA-UCEC data, and 1848 miRNAs were observed. Normalization assessment was applied analogously to the MSK test data and the TCGA-UCEC mixed-batch data using the DANA approach.

## RESULTS

In this section, we present the results of our proposed DANA approach for the MSK data and TCGA data. We show that (i) using the suggested statistical tools handling

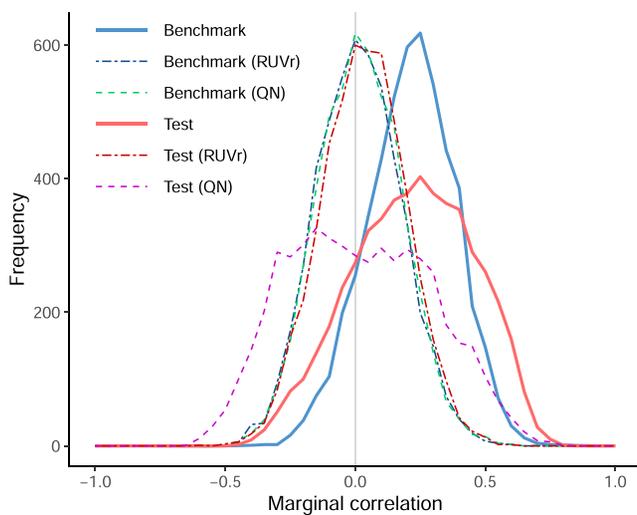


**Figure 1.** Mean count histogram (top) and mean-standard deviation plot (bottom) for the MSK test data. The ranges  $[\ell^-, u^-] = [2, 10]$  for negative controls and  $[\ell^+, \infty) = [128, \infty)$  for positive controls are indicated by blue and red vertical lines, respectively.

effects can be identified in the data, (ii) normalization is needed, (iii) the relative performance of normalization methods is data-dependent, and (iv) our approach guides the choice for a most suitable method. All results were generated in R version 4.0.2, and the code and data are available on GitHub (<https://github.com/LXQin/DANA-paper-supplementary-materials>).

### Normalization assessment for the MSK data sets

**Definition of positive and negative controls and their empirical validation.** We selected the ranges  $[\ell^+, \infty)$  and  $[\ell^-, u^-]$  for positive and negative controls, respectively, based on the observed mean count distribution and mean-standard deviation plots for the test data set (Figure 1). We set the threshold  $\ell^-$  to 2 such that (i) non-expressed miRNAs are excluded, (ii) selected miRNAs show at least mild expression, and (iii) selected miRNAs have a heterogeneous distribution of their standard deviation. We selected the thresholds  $u^- = 10$  leading to  $p^- = 102$  negative controls, and  $\ell^+ = 128$



**Figure 2.** Density curves for Pearson correlations in negative controls. Comparison between the un-normalized, RUVr-normalized, and QN-normalized MSK benchmark data, and the un-normalized, RUVr-normalized and QN-normalized MSK test data. The density curves for all normalization methods are shown in Supplementary Figures S7 and S8.

resulting in  $p^+ = 115$  positive controls. The results of the subsequent analyses are robust with respect to the selected ranges, which we show in Section 1.1 of the Supplementary Material, including Supplementary Figure S1.

*Empirical justification of the correlation measures and summary metrics.* We provide empirical evidence that supports the choice of correlation measures and the use of the numeric metrics  $m_{scr}^-$  and  $cc^+$  for quantifying the reduction of handling effects and preservation of biological signals, respectively.

For negative controls, we assess the level of their inter-marker correlation in the test data before and after normalization, and compare these with that in the benchmark data (Figure 2). On one hand, we observe that the correlations of the un-normalized data sets are not centered around zero, and the mean correlation strength is about 0.21. This was expected in both data sets due to the nature of high-throughput data, where markers in the same sample are profiled in the same experimental unit. On the other hand, the variation of correlation strengths in the test data is much higher, and, in particular, stronger positive correlations are much more abundant compared to the benchmark data. This finding accords with our assumption that handling effects manifest themselves as excessive positive correlations and reflects the higher quality of the benchmark data.

Introducing normalization to the data sets, the correlation means are shifted towards zero. While the correlation variance in the benchmark data remains approximately equal, it strongly varied for the test data depending on the applied normalization. For example, RUVr normalization leads to a strong variance reduction of inter-marker correlations in the test data and a correlation distribution similar to that of the benchmark data. On the other hand, QN, for example, centers the distribution but considerably in-

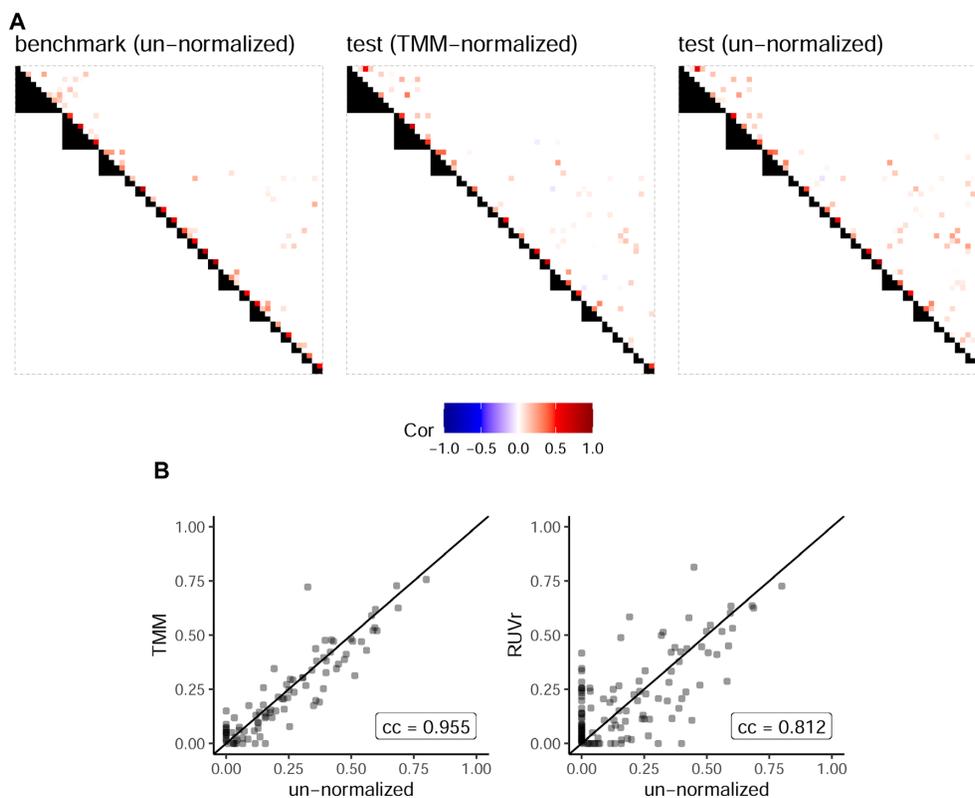
creases variance. The  $m_{scr}^-$  metric jointly captures both desired normalization effects—centering around zero and variance reduction—and, hence, offers a simple and easy-to-interpret metric to assess the reduction of handling effects in the data. For the sake of clarity, only two exemplary methods are shown in Figure 2; for all methods, see Supplementary Figures S7 and S8.

For positive controls, there is a high abundance of within-cluster correlations in the benchmark, and all of these correlations are strictly positive (Figure 3A). This aligns well with the biological evidence of co-expression of miRNAs in polycistronic clusters (20–23). Compared to the benchmark, the test data has fewer positive within-cluster correlations and more positive off-cluster correlations, both of which likely resulted from excessive handling effects. Normalization, regardless of the method, alleviated the off-cluster correlations in terms of both the number and strength. Depending on the method, normalization can either retain or reduce within-cluster correlations, which is signified by their before-versus-after-normalization concordance. The concordance is high for TMM ( $cc^+$ : 0.955) and low for RUVr ( $cc^+$ : 0.812) (Figure 3B). Hence, for the MSK test data set, TMM better preserves biological effects in positive controls than RUVr. Note that while RUVr offers the highest reduction of handling effects in negative controls, it over-corrects and fails to preserve biological signals in positive controls. Hence, it is crucial to assess normalization based on both metrics in tandem.

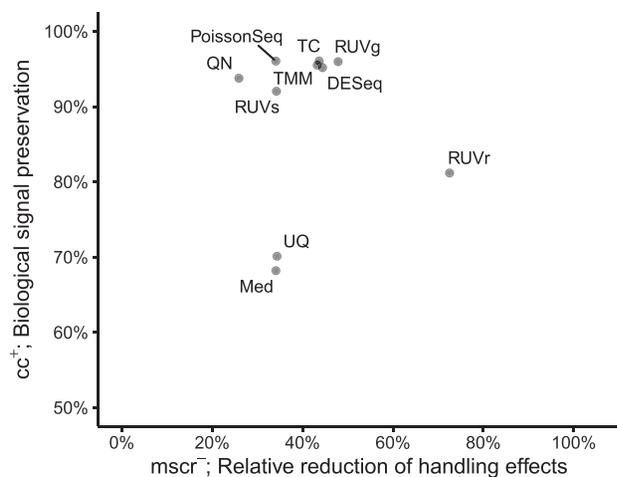
*DANA selection of a suitable normalization method for the MSK test data.* We computed the DANA correlation measures for all eight normalization methods applied to the test data (Figure 4). RUVr ( $cc^+$ : 81.2%;  $m_{scr}^-$ : 72.5%) achieved the highest reduction of handling effects, as measured by  $m_{scr}^-$ , but only moderately preserved biological signals, as measured by  $cc^+$ . RUVg (96.0%; 47.9%) offers the best compromise between a very high  $cc^+$  and a high  $m_{scr}^-$ . DESeq (95.2%; 44.4%), TMM (95.5%; 43.3%), and TC (96.1%; 43.1%) all performed well with high  $cc^+$  and relatively high  $m_{scr}^-$ , indicating a good reduction of handling effects and good preservation of biological signals. They were followed closely by PoissonSeq (96.0%; 34.1%) and RUVs (92.0%; 34.2%). While QN (93.8%; 25.9%) kept biological signals mostly intact, it reduced handling effects much less than the aforementioned methods. The method with the worst  $m_{scr}^-$  was QN (93.8%; 25.9%). Finally, UQ (70.1%; 34.4%) and Med (68.2%; 34.1%) performed poorly with a mild reduction of handling effects and the overall worst preservation of biological signals. Hence, for the MSK test data, our findings suggest choosing RUVg, TMM, TC, or DESeq for depth normalization and strongly discourage the use of Med, UQ, QN, or RUVr normalization.

The results of our normalization assessment for this data set are in accordance with our prior findings as published in (16); across both studies, the same methods were selected among the best and worst performers, respectively.

*Comparison of partial correlation estimation methods.* Computation of the  $cc^+$  metric requires the estimation of partial correlations among positive controls. While we used



**Figure 3.** (A) Partial correlation heatmaps for a subset of positive controls in the un-normalized MSK benchmark (left), TMM-normalized test (center), and un-normalized test (right) MSK data. Polycystronic clusters are indicated on the lower triangular part. (B) Partial correlation scatter plots for within-cluster correlations using TMM normalization (left) and RUVr normalization (right), each in comparison with the un-normalized test data. The concordance correlation ( $cc$ ) of within-cluster correlations before *versus* after normalization is shown in the bottom right. See Supplementary Figures S5 and S6 for the partial correlation heatmaps and scatter plots, respectively, for all methods.



**Figure 4.** DANA summary metrics  $mscr^-$  and  $cc^+$  for the MSK test data.

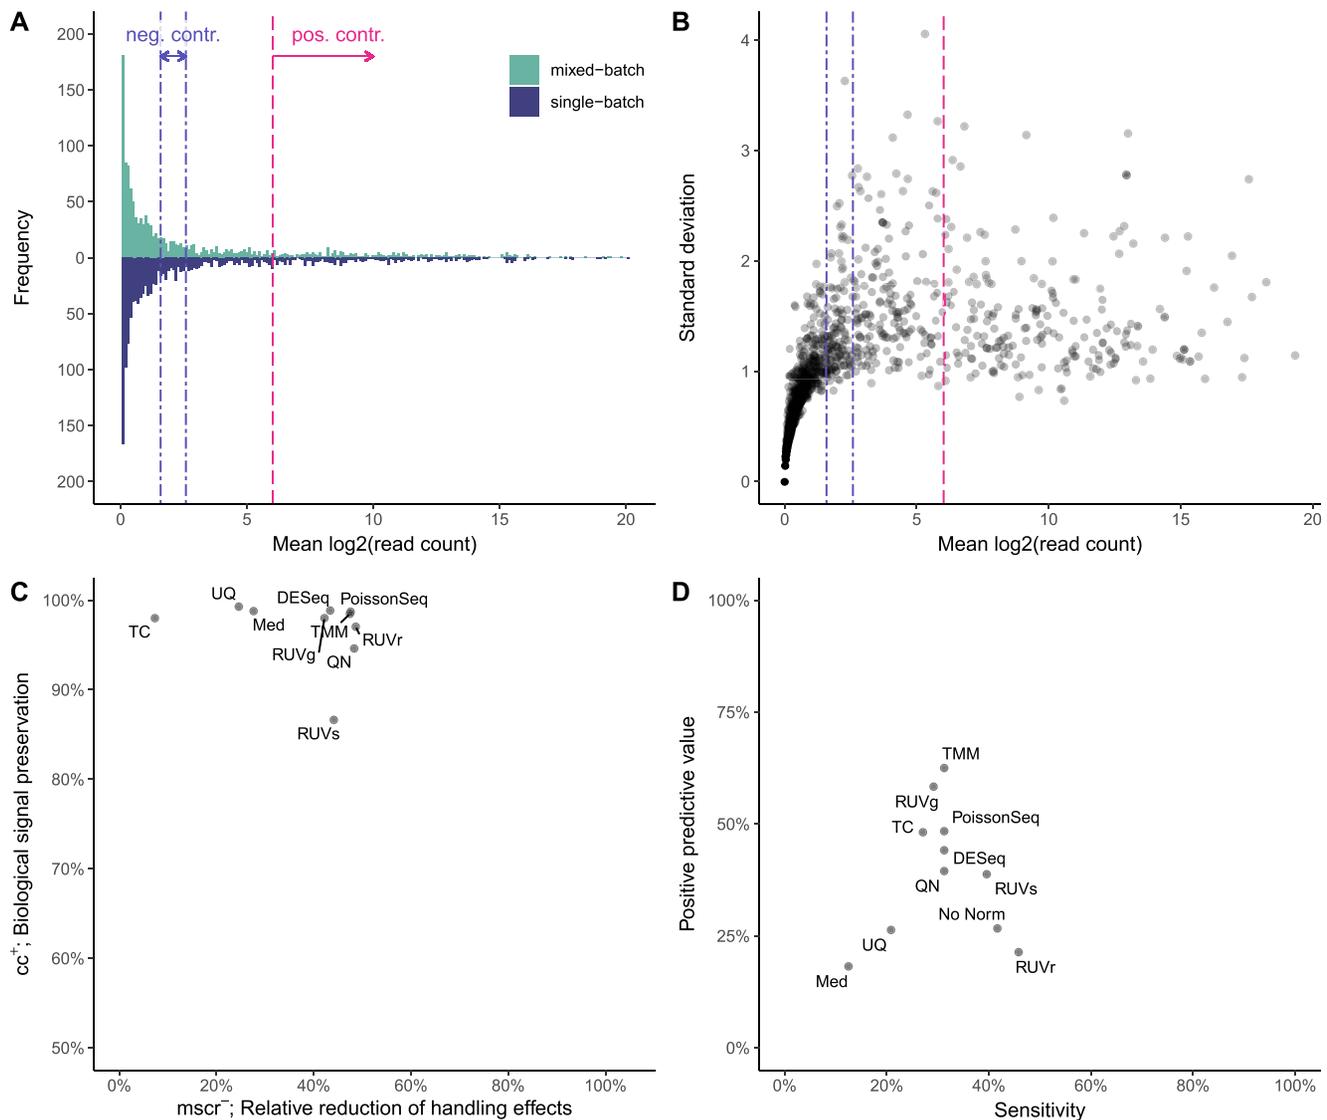
the popular neighborhood selection method (29) and calibrated its tuning parameter using Bayesian information criterion in this study, we further examined our findings for other partial correlation estimation methods. We compared the DANA results for the test data using (i) neighborhood selection with different tuning parameter calibration meth-

ods, (ii) the glasso method (30) using various tuning parameter selection methods, and (iii) the FastGGM method (31). In summary, we found that the DANA assessment is robust with respect to the chosen precision estimation method and observed a high correlation of the metric  $cc^+$  across most methods (Supplementary Figures S3 and S4). We defer a detailed comparison to Section 1.2 of the Supplementary Material.

#### Normalization assessment for the TCGA-UCEC data

For the TCGA-UCEC mixed-batch data set, we selected the ranges  $[\ell^-, u^-] = [2, 5]$  for the definition of negative controls and  $[\ell^+, \infty) = [64, \infty)$  for positive controls. This selection was based on the same reasoning we used for the MSK data and facilitated by the same graphical tools (Figures 5A, B). The number of positive and negative controls for the mixed-batch data is  $p^+ = 110$  and  $p^- = 112$ , respectively.

All considered normalization methods achieved a very high preservation of biological signals ( $cc^+ > 94\%$ ) except RUVs ( $cc^+$ : 86.6%;  $mscr^-$ : 44.1%) (Figure 5 C). However, different methods led to different reduction of handling effects. TMM (98.5%; 47.4%) and PoissonSeq (98.7%; 47.6%) offer a very high  $cc^+$  combined with a high  $mscr^-$ , whereas TC (98.0%; 7.3%), UQ (99.3%; 24.6%), and Med (98.8%; 27.6%) performed very poorly in terms of handling effects reduction. QN (94.6%; 48.3%) and

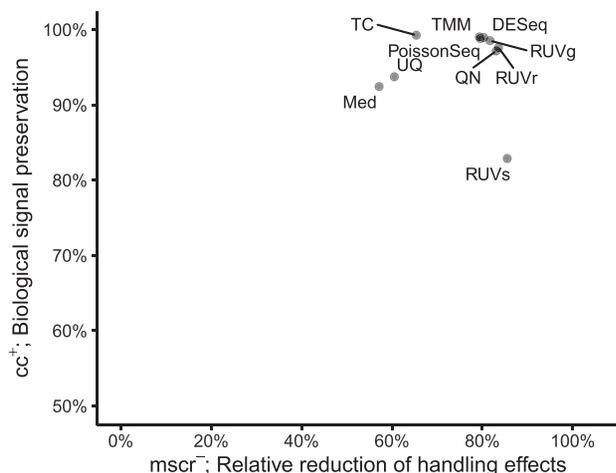


**Figure 5.** Cutoff selection and normalization assessment for the TCGA-UCEC endometrial cancer data. (A) Mean count histogram plot and (B) mean-standard deviation plot for  $\log_2$ -transformed TCGA-UCEC data. The selected ranges  $[\ell^-, u^-] = [2, 5]$  for negative controls and  $[\ell^+, \infty) = [64, \infty)$  for positive controls are indicated by blue and red vertical lines, respectively. (C) Normalization assessment using DANA in the TCGA-UCEC mixed-batch data for all normalization methods considered in this study. (D) Normalization assessment based on differential expression analysis between the subtypes END and SER. The differential expression statuses in the mixed-batch data (before and after normalization) are compared to those in the single-batch data as an assumed truth and their agreement is summarized using the positive predictive value and the sensitivity.

RUVr (97.0%; 48.7%) have the highest  $mscr^-$  but sub-optimal  $cc^+$ . Hence, TMM and PoissonSeq are most suitable for normalizing the TCGA-UCEC mixed-batch data with DESeq (98.8%; 43.4%), RUVg (98.0%; 42.2%), and RUVr as runners-up. On the other hand, we discourage using TC, UQ, Med, or RUVs normalization for this data set. The corresponding marginal correlation histograms, partial correlation heatmaps, and partial correlation scatter plots (Supplementary Figures S9–S11) further confirmed these results, see Section 2.2 of the Supplementary Material.

Among the 1848 miRNAs in the data, differential expression analysis identified 75 to be differentially expressed (DE) in the single-batch data and 48 in the un-normalized

mixed-batch data, with a significance cutoff of  $P$ -value  $< 0.01$ . We summarized the agreement of the DE statuses in the single-batch data as an assumed truth to the DE statuses in the normalized and un-normalized mixed-batch data using the positive predictive value (PPV,  $1 - \text{false discovery rate}$ ) and the sensitivity (true positive rate) (Figure 5D). Compared to the un-normalized mixed-batch (number of DE genes: 48; PPV: 27%; sensitivity: 42%), three methods decreased the PPV: Med (33; 18%; 12%), UQ (38; 26%; 21%), and RUVr (103; 21%; 46%); in addition, Med and UQ also decreased sensitivity. RUVr slightly increased the sensitivity and is the only method that increased the number of DE genes. The methods TMM (24; 62%; 31%), RUVg (24; 58%; 29%), PoissonSeq (31; 48%; 31%), DESeq (34; 44%



**Figure 6.** DANA summary metrics  $m_{scr}^-$  and  $cc^+$  for the combined TCGA-BRCA and TCGA-UCS data set.

31%), TC (27; 48%; 27%), RUVs (49; 39%; 40%), and QN (38; 39%; 31%) all increased PPV but decreased sensitivity. TMM had the overall highest PPV out of these methods and moderate sensitivity.

Taken together, in our differential expression analysis, TMM tended to outperform the other methods, while Med and UQ were the worst performers. Hence, for the TCGA-UCEC data sets, the results of our differential expression study, again, align well with the results of our DANA approach.

### Normalization assessment for the combined TCGA-BRCA and TCGA-UCS data

Lastly, we demonstrate the effectiveness of the DANA approach for the combined breast and uterine cancer data set (TCGA-BRCA and TCGA-UCS), for which no gold or silver standard is available. We selected the ranges  $[\ell^-, u^-] = [2, 5]$  and  $[\ell^+, \infty) = [100, \infty)$  based on our recommended data-driven, graphical criteria (Supplementary Figure S12). The number of negative controls is  $p^- = 91$  and the number of positive controls is  $p^+ = 116$ .

Normalization assessment through DANA (Figure 6 and Supplementary Figures S13–S15) reveals that RUVr and QN offer the best performance for the combined BRCA and UCS data set. Both methods have shown high preservation of biological signals as well as a high reduction of handling effects. As for the TCGA-UCEC data, Med, UQ, and TC again performed worse compared to all other tested normalization methods. TMM, PoissonSeq, RUVg, and DESeq showed high  $cc^+$  with intermediate  $m_{scr}^-$ . RUVs offered the highest reduction of handling effects but showed the worst biological signal preservation, highlighting the necessity of both metrics in our assessment. In summary, even though no benchmark or silver/gold standard is available, DANA could effectively assess the performance of normalization methods, and we recommend using RUVr or QN for this particular data set.

## DISCUSSION

When done appropriately, normalization can improve the accuracy and reproducibility of subsequent statistical analyses, such as the identification of disease susceptibility genes (1,39). However, it has been shown that improper normalization can lead to prominent findings that cannot be reproduced (5,40,41). Hence, the selection of a suitable normalization method is a crucial step in transcriptome sequencing data analysis. Despite numerous publications on the comparison and assessment of normalization methods, there has been no consensus on which normalization method works systematically best for which type of data. For example, the comparison of 14 studies using normalization methods for downstream differential expression analysis in (15) and the summary of the literature on normalization comparison in (42) demonstrate that no method yielded consistently high performance for different data sets. In this study, we addressed this pressing and under-studied problem by developing a data-driven approach for normalization assessment and selection in miRNA sequencing where we focused on keeping biological signals in the data intact while effectively removing unwanted variations due to heterogeneous experimental handling.

We based our approach on control markers that are defined by the data and biology, and we catered the choice of numerical measures and summary metrics to each type of control markers. Using the carefully benchmarked MSK data, we were able to empirically justify the biology-driven definition of the control markers and validate the choice of numerical measures and summary metrics. We have previously successfully used similar definitions of negative controls and positive controls for assessing the impact of normalization in the context of microarray data (15,24,42). However, we acknowledge that DANA would not be suitable for data sets in which the biological variation is expected to be zero across the samples under study for a sizable proportion of the positive control markers. Finally, we demonstrated that our approach is also robust with respect to the cutoff choices for positive and negative controls.

Across the three empirical studies we conducted in this paper, the performance of normalization methods varied greatly. A method may perform well for one data set but poorly for others. For example, based on the DANA assessment, we recommend using RUVr for the combined TCGA breast and uterine cancer data set; in stark contrast, RUVr performed poorly for the MSK sarcoma cancer data set. Hence, we confirm the necessity of a *data-driven* selection of a suitable normalization method, which previous studies are lacking. In this study, we implemented DANA—a tool that provides such data-driven guidance for miRNA data. In our three empirical studies, two methods, upper quartile and median normalization, are consistently among the worst performers based on the two DANA metrics. Furthermore, total count normalization, which is nevertheless widely used due to its simplicity, performed poorly to mediocre at best and needs to be reconsidered given our findings. Note that we do not differentiate between different types of unwanted variations in the data such as those due to different sample handling batches and those otherwise, as they tend not to be

clearly distinguished in practice. For example, TCGA uses only total count normalization and not batch effect correction for pre-processing its miRNA sequencing data, even though their data were collected in multiple batches for a disease type. However, normalization may perform poorly in presence of strong batch effects, as pointed out by a reviewer. Finally, while it is beyond the scope of this paper, the choice of sequence alignment and feature counting methods likely also affects the performance of normalization.

The application of our DANA approach does not depend on the availability of benchmark data. It applies to any set of miRNA sequencing data, such as those from TCGA. For studies with very small sample sizes, however, efforts should be made to profile the data in a single experiment with uniform handling, as DANA's metrics may not be statistically meaningful. In general, DANA can assess any normalization method that generates non-negative counts. While the assumptions of each normalization method should be checked and satisfied by the data, there are many situations in which the validity of any assumption is unknown for the given experiment, and normalization methods can perform well even if their assumptions are (partially) violated (42). To the best of our knowledge, DANA is the first approach that provides a purely data-driven assessment of normalization for miRNA sequencing data and does not depend on the availability of reliable spike-ins, housekeeping genes, or a gold standard.

## CONCLUSION

In this study, we confirmed that normalization assessment is in urgent need to objectively guide the selection of depth normalization methods for the data at hand. We developed DANA—a tool that is data-driven and biology-motivated—for guiding such selection. Our results also show that there is still a need for more effective normalization methods and support the practice of careful study design, such as uniform handling and a balanced sample-to-library-assignment, for circumventing the need for normalization. While we validated our approach using the most common downstream analysis (differential expression analysis) in this paper, we plan to test it using other analysis methods in future work. Moreover, we will extend the DANA approach towards application to data generated using other high-throughput profiling platforms. Finally, we plan to further develop our general approach for other molecules, such as mRNAs, by defining suitable sets of control markers and metrics measuring the degree of preservation of biological effects and removal of handling effects for these molecules.

## DATA AVAILABILITY

The results shown here are in part based upon data generated by the TCGA Research Network: <https://www.cancer.gov/tcga>. The MSK and TCGA data sets along with the analysis code are available on GitHub (<https://github.com/LXQin/DANA-paper-supplementary-materials>). DANA is available as an R package on GitHub (<https://github.com/LXQin/DANA>).

## SUPPLEMENTARY DATA

Supplementary Data are available at NAR Online.

## ACKNOWLEDGEMENTS

We thank the editor and the two anonymous reviewers for their insightful comments and suggestions. We also thank Joseph Christoff, Shih-Ting Huang, Mike Laszkiewicz, Nils Müller, Andy Ni, Nicole Rusk, Mahsa Taheri, Miao Wang, Yilin Wu, Qihang Yang, and Jian Zou for many helpful discussions.

## FUNDING

National Institutes of Health [CA214845 to L.X.Q. and Y.D.; HG012124 and CA008748 to L.X.Q.]; Deutsche Forschungsgemeinschaft [451920280 to Y.D. and J.L.]. Funding for open access charge: NIH [HG012124].

*Conflict of interest statement.* None declared.

## REFERENCES

- Leek, J.T., Scharpf, R.B., Bravo, H.C., Simcha, D., Langmead, B., Johnson, W.E., Geman, D., Baggerly, K. and Irizarry, R.A. (2010) Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat. Rev. Genet.*, **11**, 733–739.
- SEQC/MAQC-III Consortium (2014) A comprehensive assessment of RNA-seq accuracy, reproducibility and information content by the Sequencing Quality Control Consortium. *Nat. Biotechnol.*, **32**, 903–914.
- Risso, D., Ngai, J., Speed, T.P. and Dudoit, S. (2014) Normalization of RNA-seq data using factor analysis of control genes or samples. *Nat. Biotechnol.*, **32**, 896–902.
- Huang, H.-C., Niu, Y. and Qin, L.-X. (2015) Differential expression analysis for RNA-Seq: an overview of statistical methods and computational software. *Cancer Inform.*, **14**, 57–67.
- Rahman, M., Jackson, L.K., Johnson, W.E., Li, D.Y., Bild, A.H. and Piccolo, S.R. (2015) Alternative preprocessing of RNA-sequencing data in the Cancer Genome Atlas leads to improved analysis results. *Bioinformatics*, **31**, 3666–3672.
- Tarazona, S., Garcia-Alcalde, F., Dopazo, J., Ferrer, A. and Conesa, A. (2011) Differential expression in RNA-seq: A matter of depth. *Genome Res.*, **21**, 2213–2223.
- Mortazavi, A., Williams, B.A., McCue, K., Schaeffer, L. and Wold, B. (2008) Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nat. Methods*, **5**, 621–628.
- Robinson, M.D. and Oshlack, A. (2010) A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol.*, **11**, R25.
- Anders, S. and Huber, W. (2010) Differential expression analysis for sequence count data. *Genome Biol.*, **11**, R106.
- Dillies, M.-A., Rau, A., Aubert, J., Hennequet-Antier, C., Jeanmougin, M., Servant, N., Keime, C., Marot, G., Castel, D., Estelle, J. et al. (2013) A comprehensive evaluation of normalization methods for Illumina high-throughput RNA sequencing data analysis. *Brief. Bioinform.*, **14**, 671–683.
- Leek, J.T. (2014) svaseq: removing batch effects and other unwanted noise from sequencing data. *Nucleic Acids Res.*, **42**, 161.
- Zhang, Y., Parmigiani, G. and Johnson, W.E. (2020) ComBat-seq: batch effect adjustment for RNA-seq count data. *NAR Genom. Bioinform.*, **2**, lqaa078.
- Bullard, J.H., Purdom, E., Hansen, K.D. and Dudoit, S. (2010) Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments. *BMC Bioinform.*, **11**, 94.
- Soneson, C. and Delorenzi, M. (2013) A comparison of methods for differential expression analysis of RNA-seq data. *BMC Bioinform.*, **14**, 91.
- Li, X., Cooper, N. G.F., O'Toole, T.E. and Rouchka, E.C. (2020) Choice of library size normalization and statistical methods for

- differential gene expression analysis in balanced two-group comparisons for RNA-seq studies. *BMC Genom.*, **21**, 75.
16. Qin, L.-X., Zou, J., Shi, J., Lee, A., Mihailovic, A., Farazi, T.A., Tuschl, T. and Singer, S. (2020) Statistical assessment of depth normalization for small RNA sequencing. *JCO Clin. Cancer Inform.*, **4**, 567–582.
  17. Li, S., Labaj, P.P., Zumbo, P., Sykacek, P., Shi, W., Shi, L., Phan, J., Wu, P.-Y., Wang, M., Wang, C. *et al.* (2014) Detecting and correcting systematic variation in large-scale RNA sequencing data. *Nat. Biotechnol.*, **32**, 888–895.
  18. Ambros, V. (2004) The functions of animal microRNAs. *Nature*, **431**, 350–355.
  19. Bartel, D.P. (2004) MicroRNAs: genomics, biogenesis, mechanism, and function. *Cell*, **116**, 281–297.
  20. Baskerville, S. and Bartel, D.P. (2005) Microarray profiling of microRNAs reveals frequent coexpression with neighboring miRNAs and host genes. *RNA*, **11**, 241–247.
  21. Landgraf, P., Rusu, M., Sheridan, R., Sewer, A., Iovino, N., Aravin, A., Pfeffer, S., Rice, A., Kamphorst, A.O., Landthaler, M. *et al.* (2007) A mammalian microRNA expression atlas based on small RNA library sequencing. *Cell*, **129**, 1401–1414.
  22. Griffiths-Jones, S., Saini, H.K., van Dongen, S. and Enright, A.J. (2008) miRBase: tools for microRNA genomics. *Nucleic Acids Res.*, **36**, D154–D158.
  23. Chaulk, S.G., Ebhardt, H.A. and Fahlman, R.P. (2016) Correlations of microRNA: microRNA expression patterns reveal insights into microRNA clusters and global microRNA expression patterns. *Mol. Biosyst.*, **12**, 110–119.
  24. Qin, L.-X. and Levine, D.A. (2016) Study design and data analysis considerations for the discovery of prognostic molecular biomarkers: a case study of progression free survival in advanced serous ovarian cancer. *BMC Med. Genet.*, **9**, 27.
  25. Lu, J., Getz, G., Miska, E.A., Alvarez-Saavedra, E., Lamb, J., Peck, D., Sweet-Cordero, A., Ebert, B.L., Mak, R.H., Ferrando, A.A. *et al.* (2005) MicroRNA expression profiles classify human cancers. *Nature*, **435**, 834–838.
  26. Law, C.W., Chen, Y., Shi, W. and Smyth, G.K. (2014) voom: Precision weights unlock linear model analysis tools for RNA-seq read counts. *Genome Biol.*, **15**, R29.
  27. Zhuang, R., Simon, N. and Lederer, J. (2019) Graphical models for discrete and continuous data. arXiv doi: <https://arxiv.org/abs/1609.05551>, 15 June 2019, preprint: not peer reviewed.
  28. Lederer, J. (2021) In: *Fundamentals of High-Dimensional Statistics*. Springer International Publishing, Cham, Switzerland.
  29. Meinshausen, N. and Bühlmann, P. (2006) Variable selection and High-dimensional graphs and with the Lasso. *Ann. Stat.*, **34**, 1436–1462.
  30. Friedman, J., Hastie, T. and Tibshirani, R. (2008) Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, **9**, 432–441.
  31. Ren, Z., Sun, T., Zhang, C.-H. and Zhou, H.H. (2015) Asymptotic normality and optimalities in estimation of large Gaussian graphical models. *Ann. Stat.*, **43**, 991–1026.
  32. Lin, L.I.-K. (1989) A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, **45**, 255.
  33. Li, J., Witten, D., Johnstone, I.M. and Tibshirani, R. (2012) Normalization, testing, and false discovery rate estimation for RNA-sequencing data. *Biostatistics*, **13**, 523–538.
  34. Bolstad, B.M., Irizarry, R.A., Astrand, M. and Speed, T.P. (2003) A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics*, **19**, 185–193.
  35. Robinson, M.D., McCarthy, D.J. and Smyth, G.K. (2010) edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*, **26**, 139–140.
  36. Cancer Genome Atlas Research Network, Kandoth, C., Schultz, N., Cherniack, A.D., Akbani, R., Liu, Y., Shen, H., Robertson, A.G., Pashtan, I., Shen, R. *et al.* (2013) Integrated genomic characterization of endometrial carcinoma. *Nature*, **497**, 67–73.
  37. Cancer Genome Atlas Network (2012) Comprehensive molecular portraits of human breast tumours. *Nature*, **490**, 61–70.
  38. Cherniack, A.D., Shen, H., Walter, V., Stewart, C., Murray, B.A., Bowlby, R., Hu, X., Ling, S., Soslow, R.A., Broaddus, R.R. *et al.* (2017) Integrated molecular characterization of uterine carcinosarcoma. *Cancer Cell*, **31**, 411–423.
  39. Jaffe, A.E., Hyde, T., Kleinman, J., Weinberg, D.R., Chenoweth, J.G., McKay, R.D., Leek, J.T. and Colantuoni, C. (2015) Practical impacts of genomic data ‘cleaning’ on biological discovery using surrogate variable analysis. *BMC Bioinform.*, **16**, 372.
  40. Ransohoff, D.F. (2005) Bias as a threat to the validity of cancer molecular-marker research. *Nat. Rev. Cancer*, **5**, 142–149.
  41. Akey, J.M., Biswas, S., Leek, J.T. and Storey, J.D. (2007) On the design and analysis of gene expression studies in human populations. *Nat. Genet.*, **39**, 807–808.
  42. Evans, C., Hardin, J. and Stoebel, D.M. (2018) Selecting between-sample RNA-Seq normalization methods from the perspective of their assumptions. *Brief. Bioinform.*, **19**, 776–792.