

# SCIENTIFIC REPORTS



OPEN

## Variation-preserving normalization unveils blind spots in gene expression profiling

Carlos P. Roca<sup>1,2</sup>, Susana I. L. Gomes<sup>3</sup>, Mónica J. B. Amorim<sup>3</sup> & Janeck J. Scott-Fordsmand<sup>2</sup>

Received: 11 May 2016  
Accepted: 11 January 2017  
Published: 09 March 2017

RNA-Seq and gene expression microarrays provide comprehensive profiles of gene activity, but lack of reproducibility has hindered their application. A key challenge in the data analysis is the normalization of gene expression levels, which is currently performed following the implicit assumption that most genes are not differentially expressed. Here, we present a mathematical approach to normalization that makes no assumption of this sort. We have found that variation in gene expression is much larger than currently believed, and that it can be measured with available assays. Our results also explain, at least partially, the reproducibility problems encountered in transcriptomics studies. We expect that this improvement in detection will help efforts to realize the full potential of gene expression profiling, especially in analyses of cellular processes involving complex modulations of gene expression.

Since the discovery of DNA structure by Watson and Crick, molecular biology has progressed increasingly fast, with rapid advances in sequencing and related genomic technologies. Among these, DNA microarrays and RNA-Seq have been widely adopted to obtain gene expression profiles, by measuring the concentration of tens of thousands of mRNA molecules in single assays<sup>1–5</sup>. Despite their enormous potential<sup>6–9</sup>, problems of reproducibility and reliability<sup>10–12</sup> have discouraged their use in some areas, e.g. biomedicine<sup>13–16</sup>.

The normalization of gene expression, which is required to set a common reference level among samples<sup>17–20</sup>, has been reported to be problematic, affecting the reproducibility of results with both microarray<sup>21–23</sup> and RNA-Seq<sup>24–27</sup>. Batch effects and their influence on normalization have recently received a great deal of attention<sup>28–30</sup>, resulting in approaches aiming to remove unwanted technical variation caused by differences between batches of samples or by other sources of expression heterogeneity<sup>31–33</sup>. A different issue, however, is the underlying assumption made by the most widely used normalization methods to date, such as Median and Quantile normalization<sup>34</sup> for microarrays, or RPKM (Reads Per Kilobase per Million mapped reads)<sup>4</sup>, TMM (Trimmed Mean of M-values)<sup>35</sup>, and DESeq<sup>36</sup> normalization for RNA-Seq, which posit that all or most genes are not differentially expressed<sup>25,37–40</sup>. Although it may seem reasonable for many applications, this *lack-of-variation assumption* has not been confirmed. Moreover, results obtained with external controls<sup>37–39,41</sup> or with RT-qPCR<sup>21,24</sup> suggest that it may not be valid.

Some methods have been proposed to address this issue, based on the use of spike-ins<sup>37–39</sup>, negative control probes (SQN, Subset Quantile normalization)<sup>42</sup>, or negative control genes (RUV-2, Remove Unwanted Variation, 2-step)<sup>32</sup>. These methods use external or internal controls that are *known a priori* not to be differentially expressed<sup>43</sup>. Their applicability, however, has been limited by this requirement of a priori knowledge, which is rarely available for a sufficiently large number of controls. In addition, other methods have been proposed to address the lack-of-variation assumption by identifying a subset of non-differentially expressed genes from the assay data, such as Cross-Correlation normalization<sup>44</sup>, LVS (Least-Variant Set) normalization<sup>45</sup>, and NVAS (Nonparametric Variable Selection and Approximation) normalization<sup>46</sup>. While LVS normalization requires setting in advance a number for the fraction of genes to be considered as non-differentially expressed, with values in the range 40–60%<sup>45</sup>, Cross-Correlation and NVAS normalization are expected to degrade in performance when more than 50% of genes are differentially expressed<sup>44,46</sup>. More recently, CrossNorm has been introduced<sup>47</sup>, based on the mixture of gene expression distributions from the experimental conditions. This method, however, has been proposed for two experimental conditions, and specially for paired samples. The extension of this approach

<sup>1</sup>Department of Chemical Engineering, Universitat Rovira i Virgili, 43007 Tarragona, Spain. <sup>2</sup>Department of Bioscience, Aarhus University, 8600 Silkeborg, Denmark. <sup>3</sup>Department of Biology & CESAM, University of Aveiro, 3810-193 Aveiro, Portugal. Correspondence and requests for materials should be addressed to C.P.R. (email: carlosproca@gmail.com) or J.J.S.-F. (email: jsf@bios.au.dk)

to experimental designs with unpaired samples and more than a few experimental conditions would lead, as far as we can hypothesize, to an unmanageable size of the data matrix to process.

Thus, to clarify and overcome the limitations imposed by the lack-of-variation assumption, we have developed an approach to normalization that does not assume lack-of-variation and that is suitable to most real-world applications. Hence, we aimed to avoid the need of spike-ins, a priori knowledge of control genes, or assumptions on the number of differentially expressed genes. The analysis of several gene expression datasets using this approach confirmed that our methods reached these goals. Furthermore, our results show that assuming lack-of-variation can severely undermine the detection of gene expression variation in real assays. We have found that large numbers of differentially expressed genes, with substantial expression changes, are missed or misidentified when data are normalized with methods that assume lack-of-variation.

## Results

***E. crypticus* and Synthetic Datasets.** A large gene expression dataset was obtained from biological triplicates of *Enchytraeus crypticus* (a globally distributed soil organism used in standard ecotoxicity tests), sampled under 51 experimental conditions (42 treatments and 9 controls), involving exposure to several substances, at several concentrations and durations according to a factorial design (Supplementary Table S1). Gene expression was measured using a customized high-density oligonucleotide microarray<sup>48</sup>, resulting in a dataset with 18,339 gene probes featuring good hybridization signal in all 153 samples. Taking into account the design of the microarray<sup>48</sup>, we refer to these gene probes as *genes* in what follows.

To further explore and compare outcomes between normalization methods, two synthetic random datasets were built and analyzed. One of them was generated with identical means and variances gene-by-gene to the real *E. crypticus* dataset, and under the assumption that no gene was differentially expressed. In addition, normalization factors were applied, equal to those obtained from the real dataset. Thus, this synthetic dataset was similar to the real one, while complying by construction with the lack-of-variation assumption. The other synthetic dataset was also generated with comparable means and variances to the real dataset and with normalization factors, but in this case differential expression was added. Depending on the experimental condition, several numbers of differentially expressed genes and ratios between over- and under-expressed genes were introduced (see Methods). Together, these synthetic datasets with and without differential gene expression represent, respectively, the alternative and null hypotheses for a statistical test of differential gene expression.

**Normalization Methods.** The gene expression datasets were normalized with four methods. Two of these methods are the most widely used procedures for microarrays, namely Median (or Scale) normalization and Quantile normalization<sup>34</sup>. (Note that current methods of normalization for RNA-Seq, such as RPKM<sup>4</sup>, TMM<sup>35</sup>, and DESeq<sup>36</sup>, perform between-sample normalization by introducing a scaling per sample obtained with some form of mean or median, using all or a large set of genes. Thus their performance, in what concerns the issues addressed here, is expected to be similar to that of Median normalization for microarrays.)

The other two normalization methods were developed for this study, they being called *Median Condition-Decomposition normalization* and *Standard-Vector Condition-Decomposition normalization*, respectively MedianCD and SVCD normalization in what follows.

With the exception of Quantile normalization, all used methods apply a multiplicative factor to the expression levels of each sample, equivalent to the addition of a number in the usual  $\log_2$ -scale for gene expression levels. Solving the *normalization problem* consists of finding these correction factors. This problem can be exactly and linearly decomposed into several sub-problems: one within-condition normalization for each experimental condition and one final between-condition normalization for the condition averages (see Methods). In the within-condition normalizations, the samples (replicates) subjected to each experimental condition are normalized separately, whereas in the final between-condition normalization average levels for all conditions are normalized together. Because there are no genes with differential expression in any of the within-condition normalizations, the lack-of-variation assumption only affects the final between-condition normalization. The assumption is avoided by using, in this normalization, expression levels only from *no-variation genes*, defined as genes that show no evidence of differential expression under a statistical test. An important detail is that the within-condition normalizations ensure good estimates of the within-condition variances, which are required by the statistical test for identifying no-variation genes. This requisite also implies that a minimum of two samples is required per experimental condition. Both methods of normalization proposed here, MedianCD and SVCD normalization, follow this condition-decomposition approach.

With MedianCD normalization, all normalizations are performed with median values, as in conventional Median normalization, but only no-variation genes are employed in the between-condition step. Otherwise, if all genes were used in this final step, the resulting total normalization factors would be exactly the same as those obtained with conventional Median normalization.

For SVCD normalization, a vectorial procedure was developed to carry out each normalization step, called Standard-Vector normalization. The samples of any experimental condition, in a properly normalized dataset, must be *exchangeable*. In mathematical terms, the expression levels of each gene can be considered as an  $s$ -dimensional vector, where  $s$  is the number of samples for the experimental condition. After standardization (mean subtraction and variance scaling), these standard vectors are located in a  $(s - 2)$ -dimensional hypersphere. The exchangeability mentioned above implies that, when properly normalized, the distribution of standard vectors must be invariant with respect to permutations of the samples and must have zero expected value. These properties allow to obtain a robust estimator of the normalization factors, under fairly general assumptions that do not imply any particular distribution of gene expression (see Methods).

It is worth mentioning that the limit case when the number of samples is two ( $s = 2$ ) represents a degenerate case for Standard-Vector normalization, in which the space of standard vectors reduces to a 0-dimensional

space with only two points. In this degenerate case, Standard-Vector normalization is equivalent to global Loess normalization<sup>49,50</sup>, i.e. Loess normalization without correction for non-linearities with respect to the level of gene expression or microarray print-tips. In this sense, Standard-Vector normalization is a generalization to any number of samples of the approach underlying the different types of Loess normalization.

**Normalization Results.** Figure 1 displays the results of applying the four normalization methods to the real and two synthetic datasets. Each panel shows the interquartile range of expression levels for the 153 samples, grouped in triplicates exposed to each experimental condition. Both Median and Quantile normalization (second and third rows) yielded similar outputs for the three datasets. In contrast, MedianCD and SVCD normalization (fourth and fifth rows) detected much greater variation between conditions in the real dataset and the synthetic dataset with differential gene expression. Conventional Median normalization makes, by design, the median of all samples to be the same, while Quantile normalization makes the full distribution of gene expression of all samples to be the same. Hence, if there were differences in medians or distributions between experimental conditions, both methods would have removed them. Such variation was indeed present in the synthetic dataset with differential gene expression (Fig. 1k,n), and hence we can hypothesize the same for the real dataset (Fig. 1j,m).

**Influence of No-Variation Genes on Normalization.** To clarify how MedianCD and SVCD normalization preserved the variation between conditions, we studied the influence of the choice of no-variation genes in the final between-condition normalization. To this end, we obtained the between-condition variation as a function of the number of no-variation genes, in two families of cases. In one family, no-variation genes were chosen in decreasing order of  $p$ -values from the statistical test used to analyze variation between conditions. In the other family, genes were chosen at random. The first option was similar to the approach implemented to obtain the results presented in Fig. 1j–o, with the difference that, there, the no-variation genes were chosen automatically, by a subsequent statistical test performed on the distribution of  $p$ -values (see Methods).

For the real dataset (Fig. 2a), the random choice of genes resulted in  $n^{-1/2}$  decays ( $n$  being the number of chosen genes), followed by a plateau. The  $n^{-1/2}$  decays reflect the error in the estimation of normalization factors. Selecting the genes by decreasing  $p$ -values, however, yielded a completely different result. Up to a certain number of genes, the variance remained similar, but for larger numbers of genes the variance dropped rapidly. Figure 2a shows, therefore, that between-condition variation was removed as soon as the between-condition normalizations used genes that changed in expression level across experimental conditions. The big circles in Fig. 2a indicate the working points of the normalizations used for the results displayed in Fig. 1j,m. In fact, these points slightly underestimated the variation between conditions. Although the statistical test for identifying no-variation genes ensured that no evidence of variation was found, the expression of some selected genes varied across conditions.

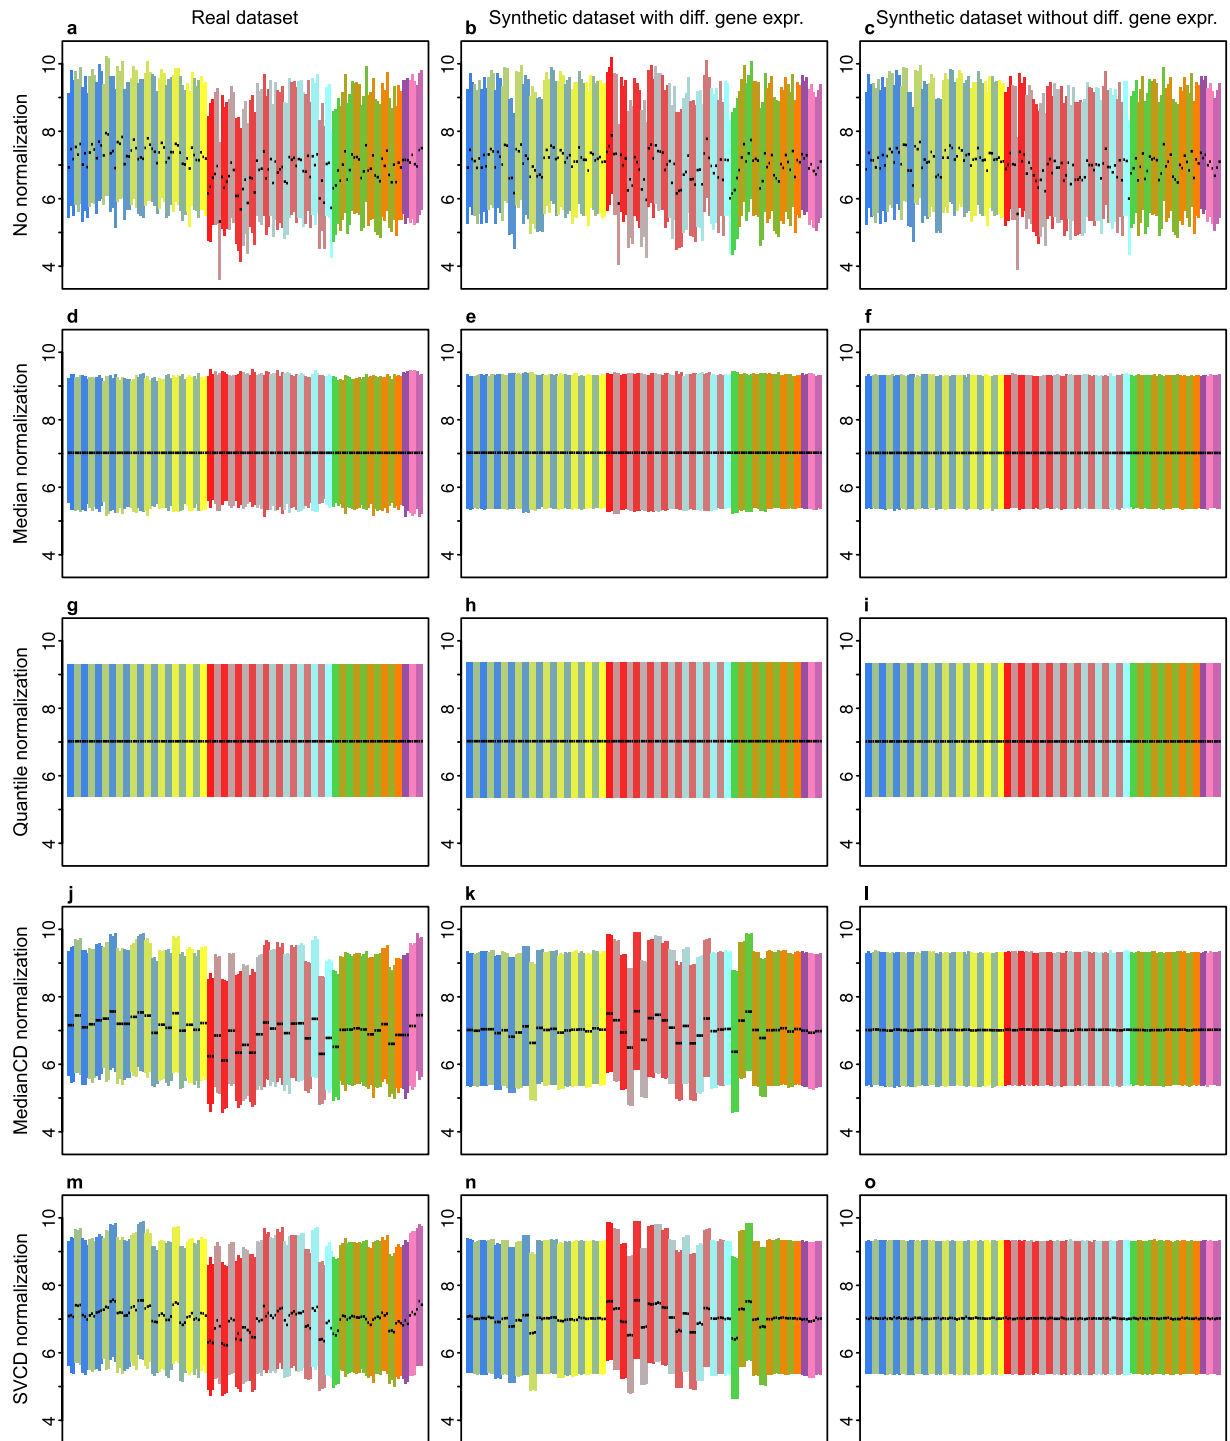
The results obtained for the synthetic dataset with differential gene expression (Fig. 2b) were qualitatively similar to those of the real dataset, but with two important differences. The amount of between-condition variation detected (by selecting no-variation genes by decreasing  $p$ -values) was smaller than with the real dataset, implying that the real dataset had larger differential gene expression. Additionally, the variation detected in the synthetic dataset had a simpler dependency on the number of genes, an indication that the differential gene expression introduced in the synthetic dataset had a simpler structure than that of the real dataset.

Figure 2c shows the results for the synthetic dataset without differential gene expression. There were no plateaus when no-variation genes were chosen randomly, only  $n^{-1/2}$  decays, and differences were small when no-variation genes were selected by decreasing  $p$ -values. Big circles show that the working points of Fig. 1l,o were selected with all available genes as no-variation genes, which is the optimum choice when there is no differential gene expression.

Overall, Fig. 2 shows that the between-condition variation displayed in Fig. 1j,k,m,n is not an artifact caused by using an exceedingly small or extremely particular set of genes in the final between-condition normalization, but that this variation originated from the datasets. The positions of the big circles in Fig. 2 highlight the good performance of the statistical approach for choosing no-variation genes in the normalizations carried out for Fig. 1j–o. Besides, the residual variation displayed by the  $n^{-1/2}$  decays implies that, as estimators of the normalization factors, SVCD normalization features smaller error than MedianCD normalization.


**Differential Gene Expression.** In what follows, we call *detected positives* the differentially expressed genes (DEGs) resulting from the statistical analyses, *treatment positives* the DEGs introduced in the synthetic dataset with differential gene expression, *true positives* the detected positives which were also treatment positives, and *false positives* the detected positives which were not treatment positives. Corresponding terms for *negatives* refer to genes which were not DEGs.

Figure 3 shows the numbers of DEGs detected in the real and synthetic datasets, for each of the 42 experimental treatments compared to the corresponding control (Supplementary Table S2), after normalizing with the four methods. For the real dataset (Fig. 3a), the number of DEGs identified after MedianCD and SVCD normalization were much larger for most treatments, in some cases by more than one order of magnitude. For the synthetic dataset with differential gene expression (Fig. 3b), results were qualitatively similar, but with less differential gene expression detected, consistently with Fig. 2a,b. The number of treatment positives can be displayed in this case (empty black down triangles, Fig. 3b), showing a better correlation, with MedianCD and SVCD normalization, between the number of treatment positives and detected positives. For the synthetic dataset without differential gene expression (Fig. 3c), no DEG was found but for one or two DEGs in two conditions. Given that the false discovery rate was controlled to be less than 5% per treatment, this is expected to happen when evaluating 42 treatments.



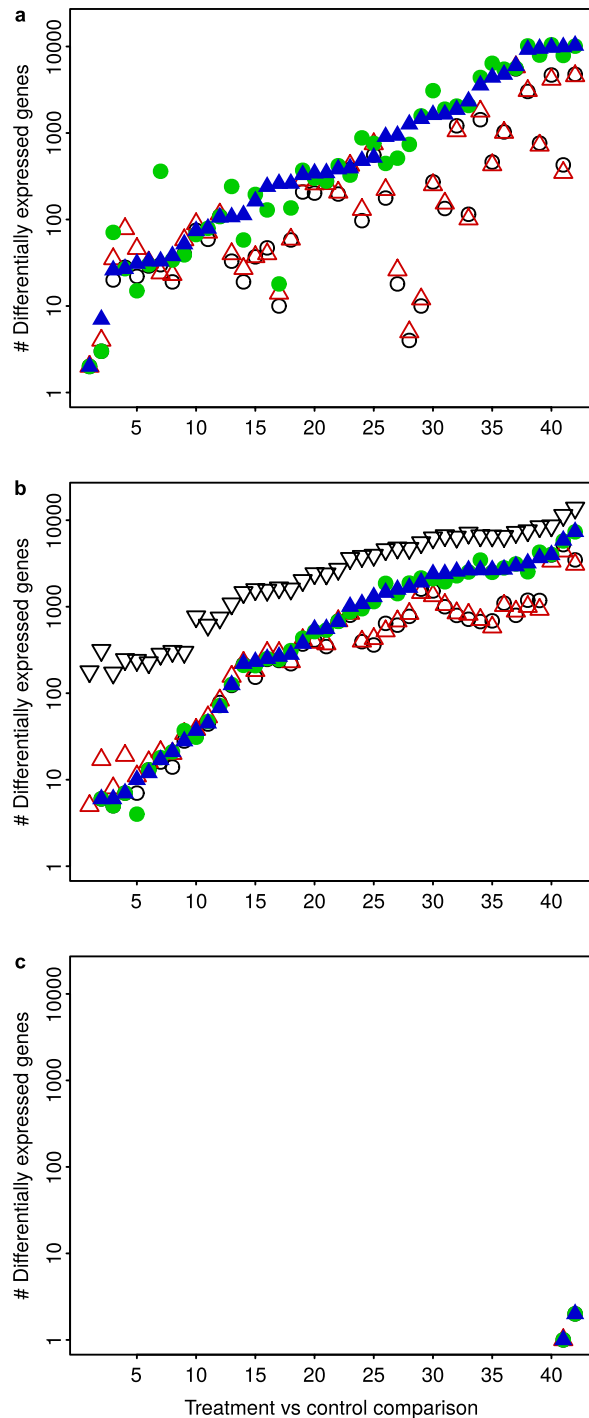
**Figure 1. MedianCD and SVCD normalization resulted in the detection of much larger between-condition variation in the datasets with differential gene expression, compared to Median and Quantile normalization.** Panels show interquartile ranges of expression levels for the 153 samples, grouped by the 51 experimental conditions (Ag, blue-yellow; Cu, red-cyan; Ni, green-orange; UV, purple; see Supplementary Table S1). Black lines indicate medians. Rows and columns correspond to normalization methods and datasets, respectively, as labeled.

Figure 3a reports, for the real dataset, statistically significant changes of gene expression, that is, changes that cannot be explained by chance. Equally important is the effect size, i.e. the scale of detected variation in DEGs, which is displayed by Fig. 4. The boxplots show absolute fold changes of expression level for all DEGs detected after applying each normalization method. MedianCD and SVCD normalization allowed to detect smaller changes of gene expression, which were otherwise missed when using Median and Quantile normalization. This differential gene expression detected with MedianCD and SVCD normalization can hardly be considered negligible, given that,

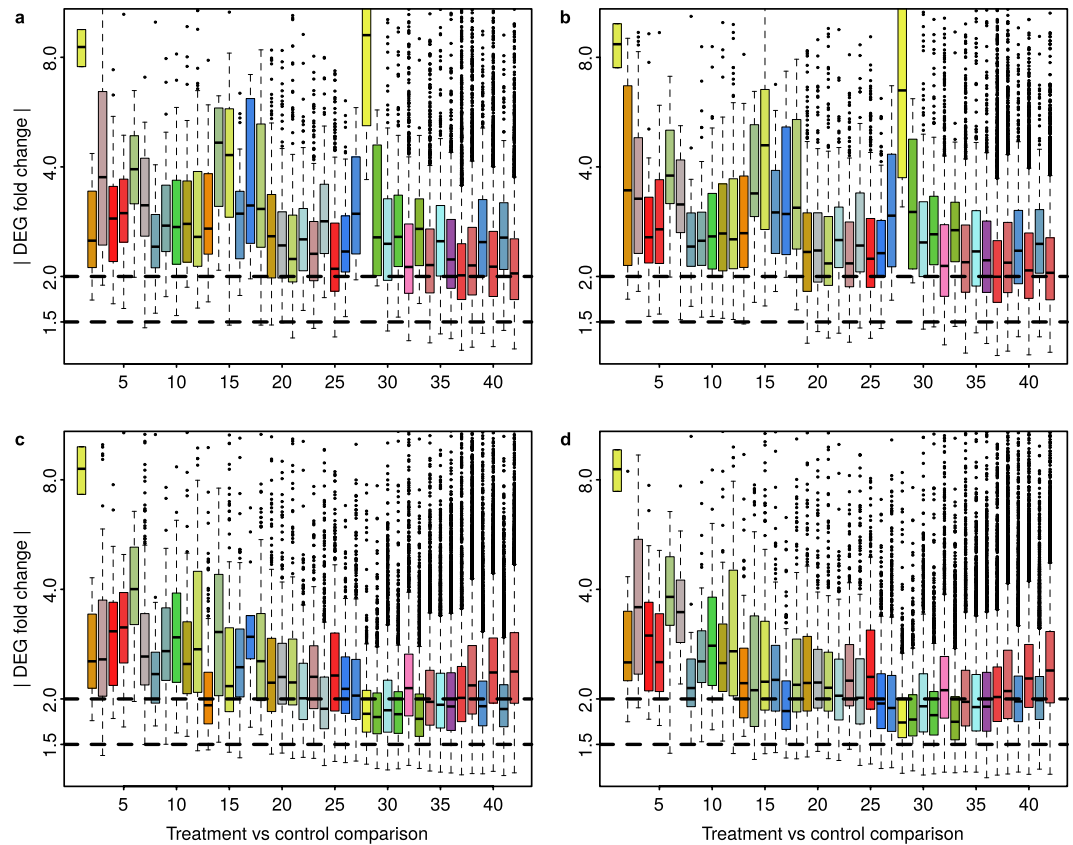


**Figure 2. The selection of genes for the final between-condition normalization in MedianCD and SVCD normalization was crucial to preserve the variation between conditions.** Panels show the detected variation as a function of the number of genes used in the between-condition normalization, for the real dataset (a), synthetic dataset with differential gene expression (b), and synthetic dataset without differential gene expression (c). Between-condition variation is represented as the standard deviation of the within-condition mean averages (averages of sample means, for all samples of the condition). See Supplementary Fig. S1 for results using within-condition median averages, with similar behavior. Each point in each panel indicates the variation obtained with one complete normalization (black circles, MedianCD normalization; blue circles, SVCD normalization). Genes were selected in two ways: randomly (empty circles) or in decreasing order of  $p$ -values from a test for detecting no-variation genes (filled circles). Big circles show the working points corresponding to the results depicted in Fig. 1j–o, which were chosen automatically. Black dashed lines show references for  $n^{-1/2}$  decays, with the same values in all panels.

for all treatments, the interquartile range of absolute fold changes was above 1.5-fold, and, for more than 28 (67%) treatments, the median absolute fold change was greater than 2-fold. Interestingly, the scale of differential gene expression detected with MedianCD and SVCD normalization in this assay is of similar magnitude to those reported by studies of global mRNA changes using external controls with microarrays and/or RNA-Seq<sup>37,39</sup>.



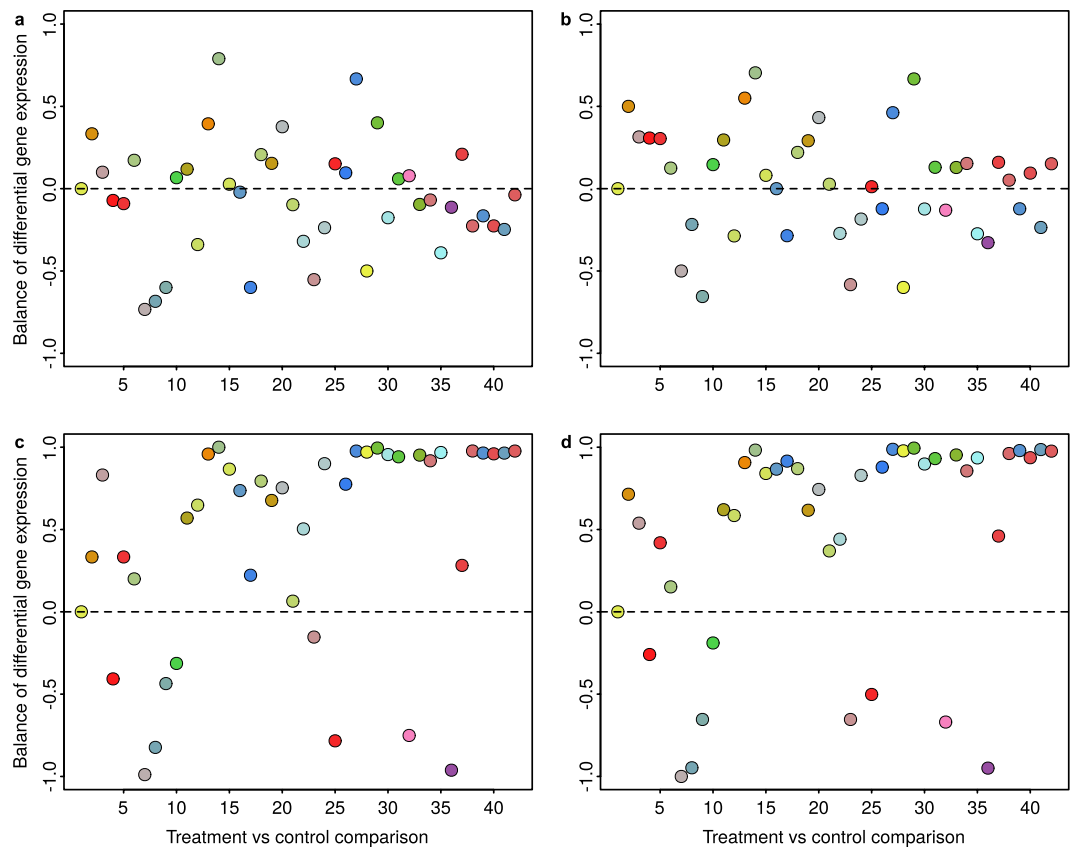
**Figure 3. MedianCD and SVCD normalization allowed to detect much larger numbers of differentially expressed genes (DEGs) in the datasets with differential gene expression.** Panels show results for the real dataset (a), synthetic dataset with differential gene expression (b), and synthetic dataset without differential gene expression (c). They display the number of DEGs for each treatment compared to the corresponding control, obtained after applying the four normalization methods (empty black circles, Median normalization; empty red up triangles, Quantile normalization; filled green circles, MedianCD normalization; filled blue up triangles, SVCD normalization). For the synthetic dataset with differential gene expression (b), the numbers of treatment positives are also shown, as empty black down triangles. In each panel, treatments are ordered according to the number of DEGs identified with SVCD normalization, increasing from left to right (see Supplementary Table 2, for real dataset). Differential gene expression was analyzed with R/Bioconductor package limma. Supplementary Fig. S2 shows results obtained with *t*-tests, qualitatively similar but with much lower detection of differential gene expression.



**Figure 4.** For the real dataset, MedianCD and SVCD normalization allowed to detect variation in gene expression of smaller magnitude than with Median and Quantile normalization. Boxplots display absolute values of DEG fold changes, for each treatment compared to the corresponding control, obtained after Median normalization (a), Quantile normalization (b), MedianCD normalization (c), and SVCD normalization (d). Boxplots are colored by treatment, with the same color code as in Fig. 1. All panels have the same order of treatments as in Fig. 3a, i.e. in increasing number of DEGs identified with SVCD normalization (Supplementary Table 2). Dashed horizontal lines indicate references of 1.5-fold and 2-fold changes.

Figure 5 displays the balance of differential gene expression, i.e. the comparison between the number of over- and under-expressed genes, for the real dataset. The quantity in the y-axis is the mean of an indicator variable  $B$ , which assigns +1 to each over-expressed DEG and -1 to each under-expressed DEG. Hence, balance of differential gene expression corresponds to  $\bar{B} = 0$ , all DEGs over-expressed to  $\bar{B} = +1$ , and, for example, 60% DEGs under-expressed to  $\bar{B} = -0.2$ . As discussed below, and as it has been reported before<sup>45,46,51,52</sup>, the balance of differential gene expression has a strong impact on the performance of normalization methods. Figure 5 shows that, regardless of the normalization method used, the unbalance of differential gene expression detected in the real dataset was substantial for most conditions. Detected unbalances were (in absolute value) larger with MedianCD and SVCD normalization, in both cases with more than 30 (71%) treatments having  $|\bar{B}| > 0.5$ , that is, more than 75% of over- or under-expressed genes. Moreover, the differences between the unbalances detected with Median and Quantile normalization, on one hand, and MedianCD and SVCD normalization, on the other, were specially notorious for the treatments with more DEGs (treatments 26–42, Fig. 3a). In those cases, Median and Quantile normalization resulted in the smallest detected unbalances, whereas MedianCD and SVCD normalization yielded the largest ones, with values near  $\bar{B} = \pm 1$  for all but two treatments.

True differential expression was known, by construction, for the synthetic dataset with differential gene expression. Thus, Fig. 6 shows for this dataset the true positive rate (ratio between true positives and treatment positives, also known as statistical power or sensitivity) and the false discovery rate (FDR, ratio between false positives and detected positives). With conditions 1 to 20, which correspond to those conditions with less than approximately 10% of treatment positives (Fig. 3b, empty black down triangles), the true positive rate was similarly low for all normalizations. Regarding the FDR, when the (total) number of detected positives was up to a few tens, variability of the FDR around the target bound at 0.05 is to be expected, given that the bound is defined over an average of repetitions of the multiple-hypothesis test. Yet, the FDR obtained after Median and Quantile normalization was higher than the 0.05 bound for most conditions. More striking, however, was the behavior for conditions 21 to 42 (more than 10% of treatment positives). The true positive rates obtained after Median and Quantile normalization were much lower than those obtained with MedianCD and SVCD normalization, while the FDR after Median and Quantile normalization was clearly over the bound at 0.05. In comparison, MedianCD



**Figure 5.** For the real dataset, detected differential gene expression was unbalanced, specially after using MedianCD and SVCD normalization and for the treatments with more DEGs (Fig. 3a). Panels show the balance of differential gene expression, for each treatment compared to the corresponding control, obtained after Median normalization (a), Quantile normalization (b), MedianCD normalization (c), and SVCD normalization (d). Each point represents the balance of differential gene expression,  $\bar{B}$  ( $\bar{B} = 0$ , same number of over- and under-expressed genes;  $\bar{B} = +1$ , all DEGs over-expressed;  $\bar{B} = -0.5$ , 75% DEGs under-expressed). Points are colored by treatment, with the same color code as in Figs 1 and 4. All panels have the same order of treatments as in Figs 3a and 4, i.e. in increasing number of DEGs identified with SVCD normalization (Supplementary Table 2). Dashed horizontal lines indicate references for balanced differential gene expression ( $\bar{B} = 0$ ).

and SVCD normalization, besides offering better sensitivity of differential gene expression, maintained the FDR consistently below the desired bound.

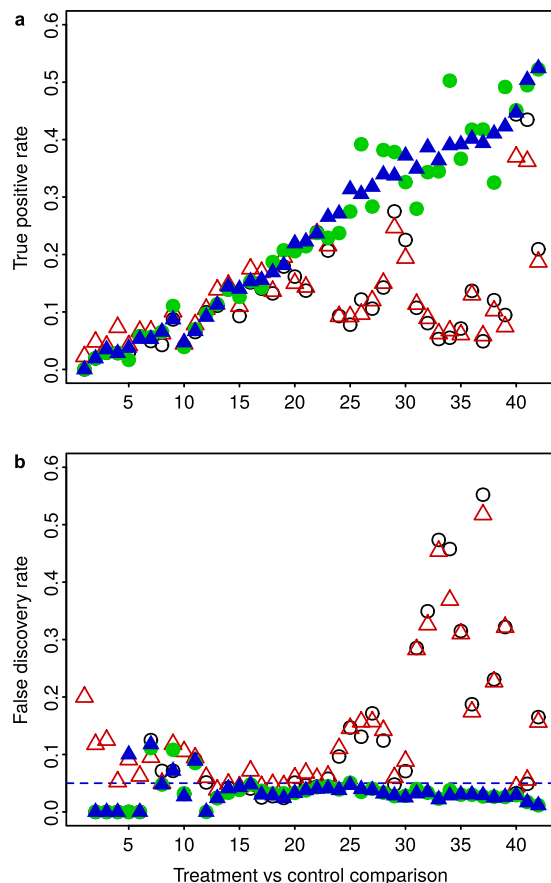
Figure 7 further explores these results, by representing the true positive rate and false discovery rate (FDR) as a function of the unbalance between over- and under-expressed genes. Figure 7 shows that the unbalance of differential gene expression was a key factor in the results obtained with Median and Quantile normalization. When most DEGs were over- or under-expressed, both the true positive rate and FDR degraded markedly after using Median or Quantile normalization. In contrast, the true positive rate and FDR were not affected by the unbalance of differential gene expression when using MedianCD or SVCD normalization.

Concerning the identification of no-variation genes, both MedianCD and SVCD normalization performed well. In the synthetic dataset without differential gene expression, both methods identified all genes as no-variation genes, which is the best possible result. In the synthetic dataset with differential gene expression, 1,834 genes (10% of a total of 18,339 genes) were, by construction, negatives across all treatments. MedianCD and SVCD normalization detected, respectively, 1,723 and 1,827 no-variation genes, among which 96.9% and 95.2% were true negatives.

**Analysis of the Golden Spike and Platinum Spike Datasets.** To provide additional evidence of the performance of MedianCD and SVCD normalization, we analyzed the Golden Spike<sup>53</sup> and Platinum Spike<sup>52</sup> datasets. Both of them are artificial real datasets, the largest ones for which true DEGs are known. Hence, they have been widely used to benchmark normalization methods<sup>45–47,52–55</sup>.

The design of the Golden Spike dataset was questioned for reasons concerning, among others, the anomalous null distribution of  $p$ -values, the lack of biological replicates, and the high concentration of spike-ins<sup>51,56,57</sup>. Nevertheless, this dataset is worth considering here because it challenges what we claim are key capabilities of our approach, that is, to correctly normalize gene expression data when many genes are differentially expressed, even with large unbalance between over- and under-expression. This dataset consists of microarray data obtained



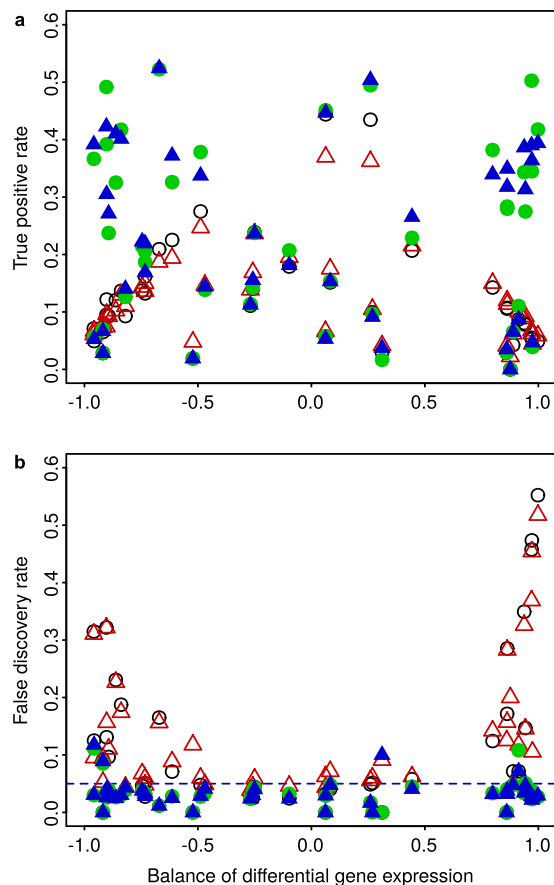


**Figure 6.** In the synthetic dataset with differential gene expression, and with more than 10% of treatment positives (Fig. 3b), Median and Quantile normalization resulted in less statistical power and uncontrolled false discovery rate. The panels display the true positive rate (a) and false discovery rate (b), for each treatment compared to the corresponding control, obtained after applying the four normalization methods (same symbols as in Fig. 3; empty black circles, Median normalization; empty red up triangles, Quantile normalization; filled green circles, MedianCD normalization; filled blue up triangles, SVCD normalization). Both panels have the same order of treatments as in Fig. 3b, i.e. in increasing number of differentially expressed genes identified with SVCD normalization. Differential gene expression was analyzed with R/Bioconductor package limma. The dashed horizontal line in (b) indicates the desired bound on the false discovery rate at 0.05.

with the Affymetrix GeneChip DrosGenome1, with two experimental conditions and three technical replicates per condition. Excluding Affymetrix internal control probes, the dataset contains a total of 13,966 gene probe sets, of which 3,876 were spiked-in, which we call *known* in what follows. Among these, 1,328 (34.3%) were over-expressed (known positives) to varying degrees between 1.1- and 4-fold, while the remaining 2,535 (65.4%) were spiked-in at the same concentration in both conditions (known negatives). (Percentages do not add up to 100% because of a very small number of probe sets with weak matching to multiple clones<sup>53</sup>).

In addition to the normalization methods used above, we included Cyclic Loess normalization<sup>49,58</sup> in this case, because it facilitates a better comparison of results with previous studies<sup>45–47,52–55</sup>. Figure 8 summarizes the results obtained for the Golden Spike dataset, by displaying Receiver Operating Characteristic (ROC) curves for the detection of differential gene expression. The upper panel shows the true positive rate (as before, ratio between true positives and treatment positives) versus the false positive rate (ratio between false positives and treatment negatives), while the lower panel shows the number of true positives versus the number of false positives. In both cases, detected and treatment positives/negatives were restricted to known genes, following previous studies<sup>54,55,57</sup>. Doing otherwise would have given an excessively dominant role to the issue of cross-hybridization in the analysis of differential gene expression<sup>54</sup>. Additionally, the analysis was performed using only probe sets with hybridization signal in all samples, with the aim of factoring out differences between normalization methods caused by the response to missing data. Results obtained without this restriction (Supplementary Fig. S3) or with *t*-tests instead of limma<sup>59</sup> analysis (Supplementary Fig. S4) were very similar to those of Fig. 8.

The comparison of ROC curves shown in Fig. 8 highlight the superior performance of MedianCD and, in particular, SVCD normalization. Dashed lines show results when the list of known negatives was given as an input to some of the normalization methods (something that cannot be done in real assays). It is remarkable that SVCD normalization featured equally well with or without this information.

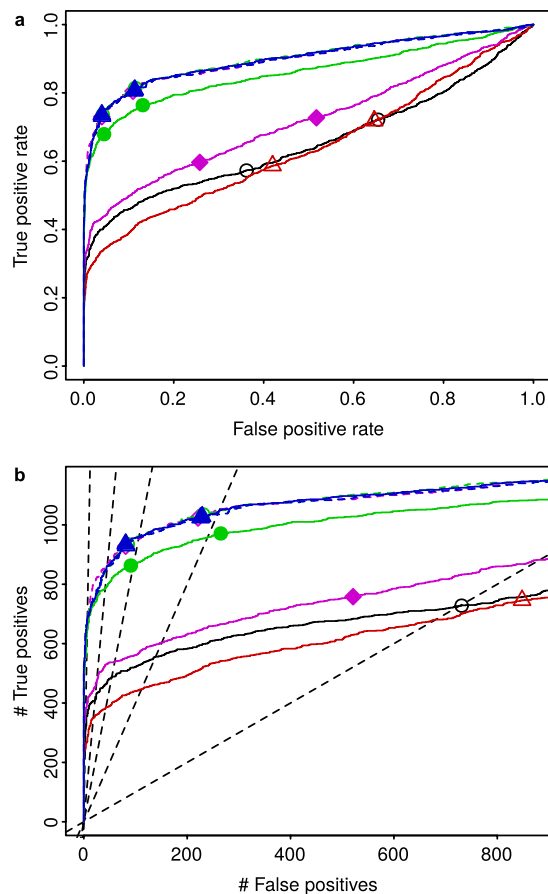


**Figure 7.** In the synthetic dataset with differential gene expression (Fig. 6), the unbalance between over- and under-expressed genes was a key factor in the lowered true positive rate and uncontrolled false discovery rate obtained after Median and Quantile normalization. Panels show the true positive rate (a) and false discovery rate (b) as a function of the balance of differential gene expression,  $\bar{B}$  ( $\bar{B} = 0$ , same number of over- and under-expressed genes;  $\bar{B} = +1$ , all DEGs over-expressed;  $\bar{B} = -0.5$ , 75% DEGs under-expressed). Each point in both panels represents the results for one treatment compared to the corresponding control, obtained after applying the four normalization methods (same symbols as in Figs 3 and 6; empty black circles, Median normalization; empty red up triangles, Quantile normalization; filled green circles, MedianCD normalization; filled blue up triangles, SVCD normalization). Differential gene expression was analyzed with R/Bioconductor package limma. The dashed horizontal line in (b) indicates the desired bound on the false discovery rate at 0.05.

Points in Fig. 8 indicate the results when controlling the false discovery rate (FDR) to be below 0.01 (left point on each curve) or 0.05 (right point). Figure 8b shows reference lines for actual FDR equal to 0.01, 0.05, 0.1, 0.2 and 0.5 (from left to right). In all cases, the FDR was not adequately controlled, although the difference between intended and actual FDR was notably smaller with MedianCD and SVCD normalization. Lack of control of the FDR in the analysis of this dataset has been previously reported<sup>53,55</sup>. It is caused by the non-uniform (hence anomalous) distribution of  $p$ -values for negative genes, which results from the analysis of differential gene expression<sup>55–57,60</sup>. It has been argued that this anomalous distribution of  $p$ -values is, in turn, a consequence of the own experimental design of the dataset, in particular the lack of biological replication and the way clone aliquots were mixed to produce each gene group with a given fold change<sup>56</sup>. Later studies have attributed this issue mostly to non-linear or intensity-dependent effects, not properly corrected in the within-sample normalization step (e.g. background correction) of the analysis pipeline<sup>52,55,57,60</sup>.

Concerning the identification of no-variation genes, both MedianCD and SVCD normalization worked correctly. MedianCD normalization identified 561 no-variation genes, of which 93.9% were known, and among which 84.1% were known negatives. SVCD normalization, in comparison, featured better detection, with 1,224 no-variation genes identified, of which 94.4% were known, and among which 90.0% were known negatives.

The design of the Platinum Spike dataset<sup>52</sup> took into account the concerns raised by the Golden Spike dataset, offering a dataset with two experimental conditions and nine (three biological  $\times$  three technical) replicates per condition, and including near 50% more spike-ins. Besides, differential gene expression was balanced, with respect to both total mRNA amount and extent of over- and under-expression. Gene expression data was obtained with Affymetrix Drosophila Genome 2.0 microarrays. Excluding Affymetrix internal control probes, the dataset contained a total of 18,769 probe sets, of which 5,587 were spiked-in, called *known* as above. Among these,

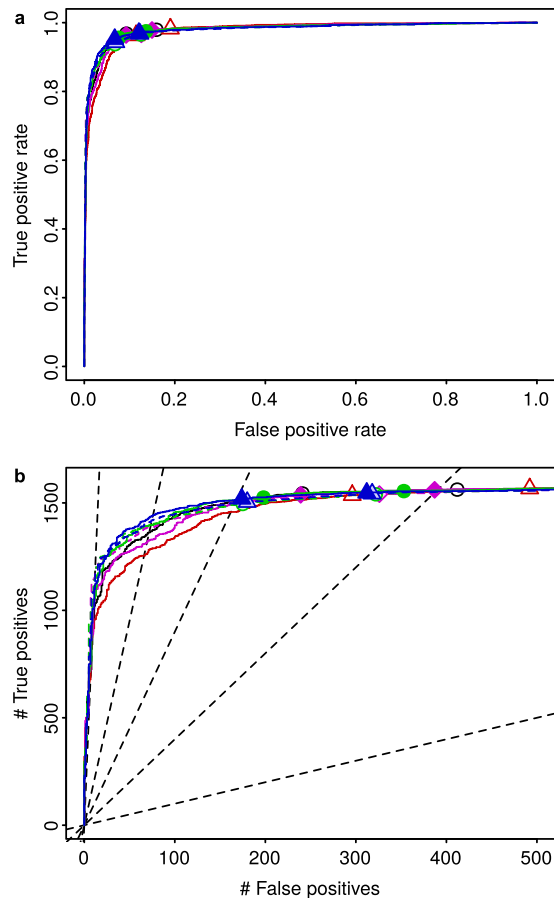


**Figure 8.** In the Golden Spike dataset, the best detection of differential gene expression was achieved after using MedianCD and, specially, SVCD normalization. Panels display ROC curves, with the true positive rate versus the false positive rate (a), or the number of true positives versus the number of false positives (b). Each curve shows the results obtained after applying the four normalization methods plus Cyclic Loess normalization (same colors and symbols as in Figs 3, 6 and 7; black curve with empty black circles, Median normalization; red curve with empty red up triangles, Quantile normalization; green curve with filled green circles, MedianCD normalization; blue curve with filled blue up triangles, SVCD normalization; magenta curve with filled magenta diamonds, Cyclic Loess normalization). Dashed curves with lightly filled symbols, overlapping the response of SVCD normalization, show results when the list of known negatives was provided to MedianCD, SVCD, and Cyclic Loess normalization. The two points per normalization method show results when controlling the false discovery rate (FDR) to be below 0.01 (left point) or 0.05 (right point). Dashed lines in (b) show references for actual FDR equal to 0.01, 0.05, 0.1, 0.2, or 0.5 (from left to right). Compared to MedianCD and SVCD normalization, the other normalization methods resulted in notably more severe degradation of the FDR.

1,940 (34.7%) were differentially expressed (known positives) to varying degrees between 1.2- and 4-fold (1,057 over-expressed, 883 under-expressed), while the remaining 3,406 (61.0%) were spiked-in at the same concentration in both conditions (known negatives).

Figure 9 shows ROC curves for the Platinum Spike dataset. As above, only known genes were considered for detected and treatment positives/negatives. Additionally, gene probes were restricted to those with signal in all samples. Results obtained without this restriction (Supplementary Fig. S5) or with *t*-tests instead of limma analysis (Supplementary Fig. S6) were again very similar. In contrast to the Golden Spike dataset (Fig. 8), the performance concerning true and false positives resulting from the different normalization methods was much more comparable. In this case, MedianCD and SVCD normalization were only marginally better. Note, however, that the FDR was again not properly controlled (Fig. 9b). Similarly to the Golden Spike dataset, and despite biological replication and a different experimental setup, obtained distributions of *p*-values for negative genes have been reported to be non-uniform<sup>52</sup>. This fact is consistent with previous arguments relating the lack of control of the FDR to a general problem concerning the correction of non-linearities in the preprocessing of microarray data<sup>52,55,57,60</sup>.

Regarding the identification of no-variation genes, MedianCD and SVCD normalization also worked correctly with this dataset. MedianCD normalization identified 2,090 no-variation genes, of which 95.4% were known, and among which 98.7% were known negatives. SVCD normalization featured slightly better, with 2,232 no-variation genes identified, of which 95.3% were known, and among which 98.3% were known negatives.



**Figure 9.** In the Platinum Spike dataset, all normalization methods resulted in similar detection of differential gene expression, with MedianCD and SVCD normalization being only marginally better. Panels display ROC curves, with the true positive rate versus the false positive rate (a), or the number of true positives versus the number of false positives (b). Each curve shows the results obtained after applying the four normalization methods plus Cyclic Loess normalization (same colors and symbols as in Figs 3, 6, 7 and 8; black curve with empty black circles, Median normalization; red curve with empty red up triangles, Quantile normalization; green curve with filled green circles, MedianCD normalization; blue curve with filled blue up triangles, SVCD normalization; magenta curve with filled magenta diamonds, Cyclic Loess normalization). Dashed curves with lightly filled symbols show results when the list of known negatives was provided to MedianCD, SVCD, and Cyclic Loess normalization. As in Fig. 8, the two points per normalization method show results when controlling the false discovery rate (FDR) to be below 0.01 (left point) or 0.05 (right point). Dashed lines in (b) show references for actual FDR equal to 0.01, 0.05, 0.1, 0.2, or 0.5 (from left to right). Compared to the Golden Spike dataset (Fig. 8), the difference between normalization methods in the resulting degradation of the FDR was smaller for this dataset.

## Discussion

The lack-of-variation assumption underlying current methods of normalization was self-fulfilling, removing variation in gene expression that was actually present. Moreover, it had negative consequences for downstream analyses, as it removed potentially important biological information and introduced errors in the detection of gene expression. The resulting decrease in statistical power or sensitivity is a handicap, which can be addressed by increasing the number of samples per experimental condition. However, degradation of the (already weak) control of the false discovery rate when using Median or Quantile normalization is a major issue for real-world applications.

The removal of variation can be understood as additive errors in the estimation of normalization factors. Considering data and errors vectorially (see Methods), the length of each vector equals, after centering and up to a constant factor, the standard deviation of the data or error. Errors of small magnitude, compared to the data variance, would only have minor effects. However, errors of similar or greater magnitude than the data variance may, depending on the vector lengths and the angle between the vectors, severely distort the observed data variance. This will, in turn, cause spurious results in the statistical analyses. Furthermore, the angles between the data and the correct normalization factors (considered as vectors) are random, given that expression data reflect biological variation while normalization factors respond to technical variation. If the assay is repeated, even with exactly the same experimental setup, the errors in the normalization factors will vary randomly, causing random spurious

results in the downstream analyses. This explains, at least partially, the lack of reproducibility found in transcriptomics studies, especially for the detection of changes in gene expression of small-to-medium magnitude (up to 2-fold), because variation of this size is more likely to be distorted by errors in the estimation of normalization factors. Accordingly, the largest differences in numbers of differentially expressed genes detected by Median and Quantile normalization, compared to MedianCD and SVCD normalization, occurred in the treatments with the smallest magnitudes of gene expression changes (Figs 3a and 4).

The variation between medians displayed in Fig. 1j,k,m,n may seem surprising, given routine expectations based on current methods (Fig. 1d,e,g,h). Nevertheless, this variation inevitably results from the unbalance between over- and under-expressed genes. As an illustration of this issue, let us consider a case with two experimental conditions, in which the average expression of a given gene is less than the distribution median under one condition, but greater than the median under the other. The variation of this gene alone will change the value of the median to the expression level of the next ranked gene. Therefore, if the number of over-expressed genes is different from the number of under-expressed genes, and enough changes cross the median boundary, then the median will substantially differ between conditions. Only when differential expression is negligible or is balanced with the respect to the median, will the median stay the same. Note that this is a related but different requirement from the number of over- and under-expressed genes being the same. This argument applies equally to any other quantile in the distribution of gene expression. The case of Quantile normalization is the least favorable, because it requires that changes of gene expression are balanced with respect to all distribution quantiles.

Compared with other normalization approaches that try to identify no-variation genes from expression data, such as Cross-Correlation<sup>44</sup>, LVS<sup>45</sup>, or NVAS<sup>46</sup> normalization, our proposal is able to work correctly with higher degrees of variation in gene expression, given that those methods are not expected to work correctly when more than 50–60% of genes vary. The reason for this difference in performance lies in that those methods use a binning strategy over the average expression between conditions (Cross-Correlation, NVAS), or need to assume an a priori fraction (usually 40–60%) of non-differentially expressed genes (LVS). When the majority of genes are differentially expressed, very few of those bins may be suitable for normalization, or the assumed fraction of non-differentially expressed genes may not hold. In contrast, our approach makes one single search in a space of  $p$ -values, and without assuming any fraction of non-differentially expressed genes. As long as there are a sufficient number of non-differentially expressed genes, of the order of several hundreds, normalization is possible, including cases with global mRNA changes or transcriptional amplification<sup>37–39,41</sup>. In general, it is a matter of comparison between the magnitude of the error in the estimation of normalization factors and the amount of biological variation. The estimation error decreases with the number of no-variation genes detected (Fig. 2), and whenever normalization error is well below biological variation, normalization between samples will be correct and beneficial for downstream analyses.

Our approach to normalization is based in four key ideas: first, decomposing the normalization by experimental conditions and normalizing separately each condition before normalizing the condition means; second, using the novel Standard-Vector normalization (or alternatively median scaling) to perform each normalization; third, identifying no-variation genes from the distribution of  $p$ -values resulting from a statistical test of variation between conditions; and fourth, employing only no-variation genes for the final between-condition normalization. These four ideas are grounded on rigorous mathematical statistics (see Methods and Supplementary Information). It is also worth noting that both Median and Standard-Vector normalization, as methods for each normalization step, are distribution-free methods; they do not assume Gaussianity or any other kind of probability distribution for the expression levels of genes. MedianCD and SVCD normalization are freely available in the R package *cdnormbio*, installable from GitHub (see Code Availability).

Previous assumptions that gene variation is rather limited could suggest that there is no need for more comprehensive normalization methods such as our proposal. In line with this, it could be argued that the amount of variation in our real (*E. crypticus*) dataset is exceptional and much larger than the variation likely to be occur in most experiments. We think that this an invalid belief. Most of the available evidence concerning widespread variation in gene expression is inadequate, because it involves circular reasoning. We have shown here that current normalization methods, used by almost all studies to date, assume no variation in gene expression between experimental conditions, and they remove it if it exists, unless it is balanced. Therefore, these methods *cannot* be used to discern the extent and balance of global variation in gene expression. Only methods that are able to normalize correctly, whatever these extent and balance are, can be trusted for this task. The fact that our methods perform well with large and unbalanced differential gene expression does not imply that they perform poorly when differential gene expression is more moderate or balanced. Our results show that this is not case. In the design of our methods, no compromise was made to achieve good performance with high variation in exchange for not so good performance with low variation. The downside of our approach lies elsewhere, in a greater algorithmic complexity and a greater demand of computing resources. Yet, we consider this a minor demand, given the capabilities of today's computers and the resources required by current high-throughput assays.

Our results have being obtained from microarray data, but similar effects are expected to be found in RNA-Seq assays. Current normalization procedures for RNA-Seq, such as RPKM<sup>4</sup>, TMM<sup>35</sup>, or DESeq<sup>36</sup>, perform between-sample normalization based on some form of global scaling and under the assumption that most genes are not differentially expressed. This makes RPKM, TMM, and DESeq normalization, in what concerns between-sample normalization and the removal or distortion of variation discussed here, similar to conventional Median normalization. An example of this issue, including results from microarray and RNA-Seq assays, has been reported in a study of the transcriptional amplification mediated by the oncogene *c-Myc*<sup>39</sup>.

Importantly, MedianCD and SVCD normalization were designed with no dependencies on any particular aspect of the technology used to globally measure gene expression, i.e. microarrays or RNA-Seq. The numbers in the input data are interpreted as steady state concentrations of mRNA molecules, in order to identify the normalization factors, and irrespectively of whether the concentrations were obtained from fluorescence intensities

of hybridized cDNA (microarrays) or from counts of fragments of cDNA (RNA-Seq). Both technologies require between-sample normalization, because in some step of the assay the total mRNA or cDNA mass in each sample must be equalized within a given range required by the experimental platform. This equalization of total mass, together with other sources of variation in the total efficiency of the assay, amounts to a factor multiplying the concentration of each mRNA species. This factor is different for each sample, and it is what between-sample normalization aims to detect and correct for. Moreover, the total mRNA mass in each sample is, in many cases, mostly determined by a few highly expressed genes, rather than an unbiased average over the total mRNA population. This makes between-sample normalization critical regarding comparisons of gene expression between different experimental conditions, as our results illustrate. It is also important to highlight that this between-sample uncertainty in the measurement of mRNA concentrations is different from other issues, such as for example non-linearities. These other problems are usually more specific to each technology, and they are the scope of within-sample normalization (e.g. background correction for microarrays and gene-length normalization for RNA-Seq), which are obviously also necessary and should be applied *before* between-sample normalization. Similarly, methods that address the influence of biological or technical confounding factors on downstream analyses, such as SVA<sup>61</sup> or PEER<sup>62</sup>, should be applied when necessary, *after* normalizing.

Finally, the significance of widespread variation in gene expression merits consideration from the viewpoint of molecular and cell biology. Established understanding about the regulation of gene expression considers it as a set of processes that generally switch on or off the expression of genes, performed mostly at transcription initiation, by the combinatorial regulation of a large number of transcription factors, and with an emphasis on gene expression programs associated with cell differentiation and development. Recent studies, however, have expanded this understanding, offering a more complex perspective on the regulation of gene expression, by identifying other rate-limiting regulation points between transcription initiation and protein translation, such as transcription elongation and termination, as well as mRNA processing, transport and degradation. Promoter-proximal pausing of RNA polymerase II (in eukaryotes)<sup>63,64</sup> and transcript elongation<sup>65</sup>, in particular, have received a great deal of attention recently, in connection with gene products involved in signal transduction pathways. These mechanisms, which seem to be highly conserved among metazoans, would allow cells to tune the expression of activated genes in response to signals concerning, for example, homeostasis, environmental stress or immune response. As an illustration, studies about the transcription amplification mediated by the oncogene *c-Myc* have uncovered that it regulates the promoter-proximal pausing of RNA polymerase II, affecting a large number of genes already activated by other regulatory mechanisms<sup>66–68</sup>. Our results for the toxicity experiment with *E. crypticus* are consistent with regulatory capabilities for broad fine-tuning of gene expression levels, far beyond what conventional methods of normalization would allow to detect. This contrast underlines that normalization methods that truly preserve variation between experimental conditions are necessary for high-throughput assays exploring genome-wide regulation of gene expression, as required by current research in molecular and cell biology.

In summary, this study proves that large numbers of genes can change in expression level across experimental conditions, and too extensively to ignore in the normalization of gene expression data. Current normalization methods for gene expression microarrays and RNA-Seq, because of a lack-of-variation assumption, likely remove and distort variation in gene expression. The normalization methods proposed here solve this problem, offering a means to investigate broad changes in gene expression that have remained hidden to date. We expect this to provide revealing insights about diverse biomolecular processes, particularly those involving substantial numbers of genes, and to assist efforts to realize the full potential of gene expression profiling.

## Methods

**Test Organism and Exposure Media.** The test species was *Enchytraeus crypticus*. Individuals were cultured in Petri dishes containing agar medium, in controlled conditions<sup>69</sup>.

For copper (Cu) exposure, a natural soil collected at Hygum, Jutland, Denmark was used<sup>69,70</sup>. For silver (Ag) and nickel (Ni) exposure, the natural standard soil LUFA 2.2 (LUFA Speyer, Germany) was used<sup>69</sup>. The exposure to ultra-violet (UV) radiation was done in ISO reconstituted water<sup>71</sup>.

**Test Chemicals.** The tested Cu forms<sup>69</sup> included copper nitrate ( $\text{Cu}(\text{NO}_3)_2 \cdot 3\text{H}_2\text{O} > 99\%$ , Sigma Aldrich), Cu nanoparticles (Cu-NPs, 20–30 nm, American Elements) and Cu nanowires (Cu-Nwires, as synthesized<sup>72</sup>).

The tested Ag forms<sup>69</sup> included silver nitrate ( $\text{AgNO}_3 > 99\%$ , Sigma Aldrich), non-coated Ag nanoparticles (Ag-NPs Non-Coated, 20–30 nm, American Elements), Polyvinylpyrrolidone (PVP)-coated Ag nanoparticles (Ag-NPs PVP-Coated, 20–30 nm, American Elements), and Ag NM300K nanoparticles (Ag NM300K, 15 nm, JRC Repository). The Ag NM300K was dispersed in 4% Polyoxyethylene Glycerol Triolaete and Polyoxyethylene (20) orbital mono-Laurat (Tween 20), thus the dispersant was tested alone as control (CTdisp).

The tested Ni forms included nickel nitrate ( $\text{Ni}(\text{NO}_3)_2 \cdot 6\text{H}_2\text{O} \geq 98.5\%$ , Fluka) and Ni nanoparticles (Ni-NPs, 20 nm, American Elements).

**Spiking Procedure.** Spiking for the Cu and Ag materials was done as previously described<sup>69</sup>. For the Ni materials, the Ni-NPs were added to the soil as powder, following the same procedure as for the Cu materials.  $\text{NiNO}_3$ , being soluble, was added to the pre-moistened soil as aqueous dispersions.

The concentrations tested were selected based on the reproduction effect concentrations  $\text{EC}_{20}$  and  $\text{EC}_{50}$ , for *E. crypticus*, within 95% of confidence intervals, being:  $\text{CuNO}_3$   $\text{EC}_{20/50} = 290/360$  mgCu/kg, Cu-NPs  $\text{EC}_{20/50} = 980/1760$  mgCu/kg, Cu-Nwires  $\text{EC}_{20/50} = 850/1610$  mgCu/kg, Cu-Field  $\text{EC}_{20/50} = 500/1400$  mgCu/kg,  $\text{AgNO}_3$   $\text{EC}_{20/50} = 45/60$  mgAg/kg, Ag-NP PVP-coated  $\text{EC}_{20/50} = 380/550$  mgAg/kg, Ag-NP Non-coated  $\text{EC}_{20/50} = 380/430$  mgAg/kg, Ag NM300K  $\text{EC}_{20/50} = 60/170$  mgAg/kg, CTdisp = 4% w/w Tween 20,  $\text{NiNO}_3$   $\text{EC}_{20/50} = 40/60$  mgNi/kg, Ni-NPs  $\text{EC}_{20/50} = 980/1760$  mgNi/kg.

Four biological replicates were performed per test condition, including controls. For Cu exposure, the control condition for all the treatments consisted of soil from a control area at Hygum site, which has a Cu background concentration of 15 mg/kg<sup>70</sup>. For Ag exposure, two control sets were performed: CT (un-spiked LUFA soil, to be the control condition for AgNO<sub>3</sub>, Ag-NPs PVP-Coated and Ag-NPs Non-Coated treatments) and CTdisp (LUFA soil spiked with the dispersant Tween 20, to be the control condition for the Ag NM300K treatments). For Ni exposure, the control consisted of un-spiked LUFA soil.

**Exposure Details.** In soil (i.e. for Cu, Ag and Ni) exposure followed the standard ERT<sup>73</sup> with adaptations as follows: twenty adults with well-developed clitellum were introduced in each test vessel, containing 20 g of moist soil (control or spiked). The organisms were exposed for three and seven days under controlled conditions of photoperiod (16:8 h light:dark) and temperature 20 ± 1°C without food. After the exposure period, the organisms were carefully removed from the soil, rinsed in deionized water and frozen in liquid nitrogen. The samples were stored at −80°C, until analysis.

For UV exposure, the test conditions<sup>71</sup> were adapted for *E. crypticus*<sup>74</sup>. The exposure was performed in 24-well plates, where each well corresponded to a replicate and contained 1 ml of ISO water and five adult organisms with clitellum. The test duration was five days, at 20 ± 1°C. The organisms were exposed to UV on a daily basis, during 15 minutes per day to two UV intensities (280–400 nm) of 1669.25 ± 50.83 and 1804.08 ± 43.10 mW/m<sup>2</sup>, corresponding to total UV doses of 7511.6 and 8118.35 J/m<sup>2</sup>, respectively. The remaining time was spent under standard laboratory illumination (16:8 h photoperiod). UV radiation was provided by an UV lamp (Spectroline XX15F/B, Spectronics Corporation, NY, USA, peak emission at 312 nm) and a cellulose acetate sheet was coupled to the lamp to cut-off UVC-range wavelengths<sup>74</sup>. Thirty two replicates per test condition (including control without UV radiation) were performed to obtain 4 biological replicates for RNA extraction, each one with 40 organisms. After the exposure period, the organisms were carefully removed from the water and frozen in liquid nitrogen. The samples were stored at −80°C, until analysis.

**RNA Extraction, Labeling and Hybridization.** RNA was extracted from each replicate, which contained a pool of 20 and 40 organisms, for soil and water exposure, respectively. Three biological replicates per test treatment (including controls) were used. Total RNA was extracted using SV Total RNA Isolation System (Promega). The quantity and purity were measured spectrophotometrically with a nanodrop (NanoDrop ND-1000 Spectrophotometer) and its quality checked by denaturing formaldehyde agarose gel electrophoresis.

500 ng of total RNA were amplified and labeled with Agilent Low Input Quick Amp Labeling Kit (Agilent Technologies, Palo Alto, CA, USA). Positive controls were added with the Agilent one-color RNA Spike-In Kit. Purification of the amplified and labeled cRNA was performed with RNeasy columns (Qiagen, Valencia, CA, USA).

The cRNA samples were hybridized on custom Gene Expression Agilent Microarrays (4 × 44k format), with a single-color design<sup>48</sup>. Hybridizations were performed using the Agilent Gene Expression Hybridization Kit and each biological replicate was individually hybridized on one array. The arrays were hybridized at 65°C with a rotation of 10 rpm, during 17 h. Afterwards, microarrays were washed using Agilent Gene Expression Wash Buffer Kit and scanned with the Agilent DNA microarray scanner G2505B.

**Data Acquisition.** Fluorescence intensity data was obtained with Agilent Feature Extraction Software v. 10.7.3.1, using recommended protocol GE1\_107\_Sep09. Quality control was done by inspecting the reports on the Agilent Spike-in control probes.

**Data Analysis.** Analyses were performed with R<sup>75</sup> v. 3.3.1, using R packages plotrix<sup>76</sup> v. 3.6.3 and RColorBrewer<sup>77</sup> v. 1.1.2, and with Bioconductor<sup>78</sup> v. 3.3 packages affy<sup>79</sup> v. 1.50.0, drosgenome1.db v.3.2.3, drosophila2.db v. 3.2.3, genefilter v. 1.54.2, and limma<sup>59</sup> v. 3.28.20. Background correction was carried out by Agilent Feature Extraction software for the real (*E. crypticus*) dataset, while the Affymetrix MAS5 algorithm, as implemented in the limma package, was used for the Golden and Platinum Spike datasets.

To ensure an optimal comparison between the different normalization methods, only gene probes with good signal quality (flag IsPosAndSignif = True) in all samples were employed for the analysis of the *E. crypticus* dataset. This implied the selection of 18,339 gene probes from a total of 43,750. For the Golden and Platinum Spike datasets, data were considered as missing when probe sets were not called present by the MAS5 algorithm.

The synthetic dataset without differential gene expression was generated gene by gene as normal variates with mean and variance equal, respectively, to the sample mean and sample variance of the expression levels for each gene, as detected from the real *E. crypticus* dataset after SVCD normalization. The synthetic dataset with differential gene expression was generated equally, except for the introduction of differences in expression averages between treatments and controls. The magnitude of the difference in averages was equal, for each differentially expressed gene (DEG), to twice the sample variance. The percentage of DEGs for each treatment was chosen randomly, in logarithmic scale, from a range between 0.9% and 90%, while ensuring that 10% of genes were not differentially expressed across the entire dataset. One third of the treatments were mostly over-expressed (for each treatment independently, the probability of a DEG being over-expressed was  $O \sim 1 - |\mathcal{N}(0,0.1^2)|$ ), one third of the treatments were mostly under-expressed ( $O \sim |\mathcal{N}(0,0.1^2)|$ ), and the remaining third had mostly balanced differential gene expression ( $O \sim \mathcal{N}(0.5,0.1^2)$ ). For both synthetic datasets, the applied normalization factors were those detected by SVCD normalization from the real *E. crypticus* dataset.

Median normalization was performed, for each sample, by subtracting the median of the distribution of expression levels, and then adding the overall median to preserve the global expression level. Quantile normalization was performed as implemented in the limma package.

The two condition-decomposition normalizations, MedianCD and SVCD, proceeded in the same way: first, independent within-condition normalization for each experimental condition, using all genes. Then, one between-condition normalization, iteratively identifying no-variation genes and normalizing until convergence of the set of no-variation genes. And finally, another between-condition normalization using only the no-variation genes detected, to calculate the between-condition normalization factors.

The criterion for convergence of MedianCD normalization was to require that the relative changes in the standard deviation of the normalization factors were less than 0.1%, or less than 10% for 10 steps in a row. In the case of SVCD normalization, convergence required that numerical errors were, compared to estimated statistical errors (see below), less than 1%, or less than 10% for 10 steps in a row. Convergence of the set of no-variation genes was achieved by intersection of the sets found during 10 additional steps under convergence conditions. These default convergence parameters were used for all the MedianCD and SVCD normalizations reported, with the exception of MedianCD with the Golden Spike dataset, which used 30% (instead of 10%) of relative change for 10 steps in a row, to reach convergence.

In SVCD normalization, the distribution of standard vectors was trimmed in each step to remove the 1% more extreme values of variance.

Differentially expressed genes were identified with limma analysis or t-tests, controlling the false discovery rate to be below 5%, independently for each comparison of treatment versus control.

The reference distributions with permutation symmetry shown in the polar plots of Supplementary Movies S1–S3 were calculated through the six possible permutations of the empirical standard vectors. The Watson  $U^2$  statistic was calculated with the two-sample test<sup>80</sup>, comparing with an equal number of samples obtained by sampling with replacement the permuted standard vectors.

**Condition Decomposition of the Normalization Problem.** In a gene expression dataset with  $g$  genes,  $c$  experimental conditions and  $n$  samples per condition, the *observed* expression levels of gene  $j$  in condition  $k$ ,  $\mathbf{y}_j^{(k)} = (y_{1j}^{(k)}, \dots, y_{nj}^{(k)})'$ , can be expressed in  $\log_2$ -scale as

$$\mathbf{y}_j^{(k)} = \mathbf{x}_j^{(k)} + \mathbf{a}^{(k)}, \quad (1)$$

where  $\mathbf{x}_j^{(k)}$  is the vector of *true* gene expression levels and  $\mathbf{a}^{(k)}$  is the vector of normalization factors.

Given a sample vector  $\mathbf{x}$ , the mean vector is  $\bar{\mathbf{x}} = \bar{\mathbf{x}}\mathbf{1}$ , and the residual vector is  $\tilde{\mathbf{x}} = \mathbf{x} - \bar{\mathbf{x}}$ . Then, (1) can be linearly decomposed into

$$\bar{\mathbf{y}}_j^{(k)} = \bar{\mathbf{x}}_j^{(k)} + \bar{\mathbf{a}}^{(k)}, \quad (2)$$

$$\tilde{\mathbf{y}}_j^{(k)} = \tilde{\mathbf{x}}_j^{(k)} + \tilde{\mathbf{a}}^{(k)}. \quad (3)$$

Equations (3) define the within-condition normalizations for each condition  $k$ . The scalar values in (2) are used to obtain the equations on condition means,

$$\bar{\mathbf{y}}_j^* = \bar{\mathbf{x}}_j^* + \bar{\mathbf{a}}^*, \quad (4)$$

$$\tilde{\mathbf{y}}_j^* = \tilde{\mathbf{x}}_j^* + \tilde{\mathbf{a}}^*. \quad (5)$$

The between-condition normalization is defined by (5). Equations (4) reduce to a single number, which is irrelevant to the normalization. The complete solution for each condition is obtained with  $\mathbf{a}^{(k)} = \bar{\mathbf{a}}^{(k)} + \tilde{\mathbf{a}}^{(k)}$ .

For full details about this condition-decomposition approach, see Supplementary Mathematical Methods in the Supplementary Information.

**Standard-Vector Normalization.** The  $n$  samples of gene  $j$  in a given condition can be modeled with the random vectors  $\mathbf{X}_j, \mathbf{Y}_j \in \mathbb{R}^n$ . Again,  $\mathbf{Y}_j = \mathbf{X}_j + \mathbf{a}$ , where  $\mathbf{a}$  is a fixed vector of normalization factors. It can be proved under fairly general assumptions (see Supplementary Information), that the true standard vectors have zero expected value

$$\mathbb{E} \left[ \sqrt{n-1} \frac{\tilde{\mathbf{X}}_j}{\|\tilde{\mathbf{X}}_j\|} \right] = \mathbf{0}, \quad (6)$$

whereas the observed standard vectors verify, as long as  $\mathbf{a} \neq \mathbf{0}$ ,

$$0 < \mathbb{E} \left[ \sqrt{n-1} \frac{\tilde{\mathbf{Y}}_j}{\|\tilde{\mathbf{Y}}_j\|} \right] \frac{\tilde{\mathbf{a}}}{\|\tilde{\mathbf{a}}\|} < \mathbb{E} \left[ \sqrt{n-1} \frac{\mathbf{1}}{\|\tilde{\mathbf{Y}}_j\|} \right] \|\tilde{\mathbf{a}}\|. \quad (7)$$

This motivates the following iterative procedure to solve (3) and (5) (*standard-vector normalization*):

$$\hat{\mathbf{y}}_j^{(0)} = \tilde{\mathbf{y}}_j, \quad (8)$$



$$\hat{\mathbf{y}}_j^{(t)} = \hat{\mathbf{y}}_j^{(t-1)} - \hat{\mathbf{b}}^{(t-1)}, \quad \text{for } t \geq 1, \quad (9)$$

$$\hat{\mathbf{b}}^{(t)} = \frac{\sum_{j=1}^g \frac{\hat{\mathbf{y}}_j^{(t)}}{\|\hat{\mathbf{y}}_j^{(t)}\|}}{\sum_{j=1}^g \frac{1}{\|\hat{\mathbf{y}}_j^{(t)}\|}}, \quad \text{for } t \geq 0. \quad (10)$$

At convergence,  $\lim_{t \rightarrow \infty} \hat{\mathbf{b}}^{(t)} = \mathbf{0}$ , which implies  $\lim_{t \rightarrow \infty} \hat{\mathbf{y}}_j^{(t)} = \tilde{\mathbf{x}}_j$  and  $\sum_{t=0}^{\infty} \hat{\mathbf{b}}^{(t)} = \tilde{\mathbf{a}}$ . Convergence is faster the more symmetric the empirical distribution of  $\tilde{\mathbf{x}}_j / \|\tilde{\mathbf{x}}_j\|$  is on the unit  $(n-2)$ -sphere. Convergence is optimal with spherically symmetric distributions, such as the Gaussian distribution, because in that case

$$\mathbb{E} \left( \frac{\tilde{\mathbf{Y}}_j}{\|\tilde{\mathbf{Y}}_j\|} \right) = \lambda \tilde{\mathbf{a}}, \quad \text{with } 0 < \lambda < \mathbb{E} \left( \frac{1}{\|\tilde{\mathbf{Y}}_j\|} \right). \quad (11)$$

Assuming no dependencies between genes, an approximation of the statistical error at step  $t$  can be obtained with

$$\mathbb{E} \left( \|\hat{\mathbf{b}}^{(t)}\| \right) \approx \frac{\sqrt{g}}{\sum_{j=1}^g \frac{1}{\|\hat{\mathbf{y}}_j^{(t)}\|}}. \quad (12)$$

This statistical error was compared with the numerical error to assess convergence.

See Supplementary Mathematical Methods in the Supplementary Information for full details about this algorithm. See also Supplementary Movies S1–S3 for normalization examples.

**Identification of No-Variation Genes.** No-variation genes were identified with one-sided Kolmogorov-Smirnov tests, as goodness-of-fit tests against the uniform distribution, carried out on a distribution of  $p$ -values. These  $p$ -values were obtained from ANOVA tests on the expression levels of genes, grouped by experimental condition. The KS test was rejected at  $\alpha = 0.001$ .

See Supplementary Mathematical Methods in the Supplementary Information for more details about this approach to identify no-variation genes. See also Supplementary Movies S4–S6 for examples of use.

**Data Availability.** MIAME-compliant microarray data were submitted to the Gene Expression Omnibus (GEO) repository at the NCBI website (platform: GPL20310; series: GSE69746, GSE69792, GSE69793 and GSE69794). The novel normalization methods were implemented as R functions. This code, together with R scripts that generate the synthetic datasets and reproduce all reported results starting from the raw microarray data, are available at the GitHub repository <https://github.com/carlosproca/geneexpr-norm-paper> and in the Supplementary Data File `gene-expr-norm.zip`. Additionally, MedianCD and SVCD normalization are available via the R package `cdnormbio`, installable from the GitHub repository <https://github.com/carlosproca/cdnormbio>.

## References

- Schena, M., Shalon, D., Davis, R. W. & Brown, P. O. Quantitative monitoring of gene expression patterns with a complementary dna microarray. *Science* **270**, 467–470 (1995).
- Lockhart, D. J. *et al.* Expression monitoring by hybridization to high-density oligonucleotide arrays. *Nat Biotechnol* **14**, 1675–1680 (1996).
- Duggan, D. J., Bittner, M., Chen, Y., Meltzer, P. & Trent, J. M. Expression profiling using cdna microarrays. *Nat Genet* **21**, 10–14 (1999).
- Mortazavi, A., Williams, B. A., McCue, K., Schaeffer, L. & Wold, B. Mapping and quantifying mammalian transcriptomes by rna-seq. *Nat Methods* **5**, 621–628 (2008).
- Wang, Z., Gerstein, M. & Snyder, M. Rna-seq: a revolutionary tool for transcriptomics. *Nat Rev Genet* **10**, 57–63 (2009).
- Golub, T. R. *et al.* Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science* **286**, 531–537 (1999).
- van 't Veer, L. J. *et al.* Gene expression profiling predicts clinical outcome of breast cancer. *Nature* **415**, 530–536 (2002).
- Ivanova, N. B. *et al.* A stem cell molecular signature. *Science* **298**, 601–604 (2002).
- Chi, J.-T. *et al.* Endothelial cell diversity revealed by global expression profiling. *Proc Natl Acad Sci USA* **100**, 10623–10628 (2003).
- Tan, P. K. *et al.* Evaluation of gene expression measurements from commercial microarray platforms. *Nucleic Acids Res* **31**, 5676–5684 (2003).
- Frantz, S. An array of problems. *Nat Rev Drug Discov* **4**, 362–363 (2005).
- Couzin, J. Genomics. microarray data reproduced, but some concerns remain. *Science* **313**, 1559 (2006).
- Michiels, S., Koscielny, S. & Hill, C. Prediction of cancer outcome with microarrays: A multiple random validation strategy. *Lancet* **365**, 488–492 (2005).
- Weigelt, B. & Reis-Filho, J. S. Molecular profiling currently offers no more than tumour morphology and basic immunohistochemistry. *Breast Cancer Res* **12** Suppl. 4, S5 (2010).
- Brettingham-Moore, K. H., Duong, C. P., Heriot, A. G., Thomas, R. J. S. & Phillips, W. A. Using gene expression profiling to predict response and prognosis in gastrointestinal cancers—the promise and the perils. *Ann of Surg Oncol* **18**, 1484–1491 (2011).

16. Boutros, P. C. The path to routine use of genomic biomarkers in the cancer clinic. *Genome Res.* **25**, 1508–13 (2015).
17. Irizarry, R. A. *et al.* Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics* **4**, 249–264 (2003).
18. Tarca, A. L., Romero, R. & Draghici, S. Analysis of microarray experiments of gene expression profiling. *Am. J. Obstet. Gynecol.* **195**, 373–388 (2006).
19. Garber, M., Grabherr, M. G., Guttman, M. & Trapnell, C. Computational methods for transcriptome annotation and quantification using rna-seq. *Nat Methods* **8**, 469–477 (2011).
20. Conesa, A. *et al.* A survey of best practices for RNA-seq data analysis. *Genome Biol.* **17**, 13 (2016).
21. Shi, L. *et al.* The microarray quality control (maqc) project shows inter- and intraplatform reproducibility of gene expression measurements. *Nat Biotechnol* **24**, 1151–1161 (2006).
22. Shippy, R. *et al.* Using rna sample titrations to assess microarray platform performance and normalization techniques. *Nat Biotechnol* **24**, 1123–1131 (2006).
23. Draghici, S., Khatri, P., Eklund, A. C. & Szallasi, Z. Reliability and reproducibility issues in dna microarray measurements. *Trends Genet* **22**, 101–109 (2006).
24. Bullard, J. H., Purdom, E., Hansen, K. D. & Dudoit, S. Evaluation of statistical methods for normalization and differential expression in mrna-seq experiments. *BMC Bioinformatics* **11**, 94 (2010).
25. Dillies, M.-A. *et al.* A comprehensive evaluation of normalization methods for Illumina high-throughput rna sequencing data analysis. *Brief Bioinform* **14**, 671–683 (2013).
26. Su, Z. *et al.* A comprehensive assessment of RNA-seq accuracy, reproducibility and information content by the Sequencing Quality Control Consortium. *Nat Biotechnol* **32**, 903–914 (2014).
27. Lin, Y. *et al.* Comparison of normalization and differential expression analyses using RNA-Seq data from 726 individual *Drosophila melanogaster*. *BMC Genomics* **17**, 28 (2016).
28. Leek, J. T. *et al.* Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat Rev Genet* **11**, 733–739 (2010).
29. Reese, S. E. *et al.* A new statistic for identifying batch effects in high-throughput genomic data that uses guided principal component analysis. *Bioinformatics* **29**, 2877–83 (2013).
30. Li, S. *et al.* Detecting and correcting systematic variation in large-scale rna sequencing data. *Nat Biotechnol* **32**, 888–895 (2014).
31. Listgarten, J., Kadie, C., Schadt, E. E. & Heckerman, D. Correction for hidden confounders in the genetic analysis of gene expression. *Proc Natl Acad Sci USA* **107**, 16465–70 (2010).
32. Gagnon-Bartsch, J. A. & Speed, T. P. Using control genes to correct for unwanted variation in microarray data. *Biostatistics* **13**, 539–52 (2012).
33. Risso, D., Ngai, J., Speed, T. P. & Dudoit, S. Normalization of rna-seq data using factor analysis of control genes or samples. *Nat Biotechnol* **32**, 896–902 (2014).
34. Bolstad, B. M., Irizarry, R. A., Astrand, M. & Speed, T. P. A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics* **19**, 185–193 (2003).
35. Robinson, M. D. & Oshlack, A. A scaling normalization method for differential expression analysis of rna-seq data. *Genome Biol* **11**, R25 (2010).
36. Anders, S. & Huber, W. Differential expression analysis for sequence count data. *Genome Biol.* **11**, R106 (2010).
37. van de Peppel, J. *et al.* Monitoring global messenger RNA changes in externally controlled microarray experiments. *EMBO Rep.* **4**, 387–393 (2003).
38. Hannah, M. A., Redestig, H., Leisse, A. & Willmitzer, L. Global mRNA changes in microarray experiments. *Nat. Biotechnol.* **26**, 741–742 (2008).
39. Lovén, J. *et al.* Revisiting global gene expression analysis. *Cell* **151**, 476–482 (2012).
40. Hicks, S. C. & Irizarry, R. A. quantro: a data-driven approach to guide the choice of an appropriate normalization method. *Genome Biol* **16**, 117 (2015).
41. Hannah, M. A., Heyer, A. G. & Hinch, D. K. A global survey of gene regulation during cold acclimation in *Arabidopsis thaliana*. *PLoS Genet.* **1**, e26 (2005).
42. Wu, Z. & Aryee, M. J. Subset quantile normalization using negative control features. *J Comput Biol* **17**, 1385–1395 (2010).
43. Lippa, K. A., Duerwer, D. L., Salit, M. L., Game, L. & Causton, H. C. Exploring the use of internal and external controls for assessing microarray technical performance. *BMC Res Notes* **3**, 349 (2010).
44. Chua, S.-W. *et al.* A novel normalization method for effective removal of systematic variation in microarray data. *Nucleic Acids Res.* **34**, e38 (2006).
45. Calza, S., Valentini, D. & Pawitan, Y. Normalization of oligonucleotide arrays based on the least-variant set of genes. *BMC Bioinformatics* **9**, 140 (2008).
46. Ni, T. T., Lemon, W. J., Shyr, Y. & Zhong, T. P. Use of normalization methods for analysis of microarrays containing a high degree of gene effects. *BMC Bioinformatics* **9**, 505 (2008).
47. Cheng, L., Lo, L.-Y., Tang, N. L. S., Wang, D. & Leung, K.-S. CrossNorm: a novel normalization strategy for microarray data in cancers. *Sci. Rep.* **6**, 18898 (2016).
48. Castro-Ferreira, M. P. *et al.* Transcriptome assembly and microarray construction for *enchytraeus crypticus*, a model oligochaete to assess stress response mechanisms derived from soil conditions. *BMC Genomics* **15**, 302 (2014).
49. Yang, Y. H. *et al.* Normalization for cDNA microarray data: a robust composite method addressing single and multiple slide systematic variation. *Nucleic Acids Res.* **30**, e15 (2002).
50. Smyth, G. K. & Speed, T. Normalization of cDNA microarray data. *Methods* **31**, 265–273 (2003).
51. Irizarry, R. A., Cope, L. M. & Wu, Z. Feature-level exploration of a published Affymetrix GeneChip control dataset. *Genome Biol.* **7**, 404 (2006).
52. Zhu, Q., Miecznikowski, J. C. & Halfon, M. S. Preferred analysis methods for Affymetrix GeneChips. II. An expanded, balanced, wholly-defined spike-in dataset. *BMC Bioinformatics* **11**, 285 (2010).
53. Choe, S. E., Boutros, M., Michelson, A. M., Church, G. M. & Halfon, M. S. Preferred analysis methods for Affymetrix GeneChips revealed by a wholly defined control dataset. *Genome Biol.* **6**, R16 (2005).
54. Schuster, E. F., Blanc, E., Partridge, L. & Thornton, J. M. Estimation and correction of non-specific binding in a large-scale spike-in experiment. *Genome Biol.* **8**, R126 (2007).
55. Pearson, R. D. A comprehensive re-analysis of the Golden Spike data: towards a benchmark for differential expression methods. *BMC Bioinformatics* **9**, 164 (2008).
56. Dabney, A. R. & Storey, J. D. A reanalysis of a published Affymetrix GeneChip control dataset. *Genome Biol.* **7**, 401 (2006).
57. Gaile, D. P. & Miecznikowski, J. C. Putative null distributions corresponding to tests of differential expression in the Golden Spike dataset are intensity dependent. *BMC Genomics* **8**, 105 (2007).
58. Ballman, K. V., Grill, D. E., Oberg, A. L. & Therneau, T. M. Faster cyclic loess: normalizing RNA arrays via linear models. *Bioinformatics* **20**, 2778–2786 (2004).
59. Ritchie, M. E. *et al.* *limma* powers differential expression analyses for rna-sequencing and microarray studies. *Nucleic Acids Res* **43**, e47 (2015).

60. Fodor, A. A., Tickle, T. L. & Richardson, C. Towards the uniform distribution of null *P* values on Affymetrix microarrays. *Genome Biol.* **8**, R69 (2007).
61. Leek, J. T. & Storey, J. D. Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genet* **3**, 1724–35 (2007).
62. Stegle, O., Parts, L., Durbin, R. & Winn, J. A bayesian framework to account for complex non-genetic factors in gene expression levels greatly increases power in eQTL studies. *PLoS Comput Biol* **6**, e1000770 (2010).
63. Core, L. J., Waterfall, J. J. & Lis, J. T. Nascent RNA sequencing reveals widespread pausing and divergent initiation at human promoters. *Science* **322**, 1845–8 (2008).
64. Adelman, K. & Lis, J. T. Promoter-proximal pausing of RNA polymerase II: emerging roles in metazoans. *Nat. Rev. Genet.* **13**, 720–31 (2012).
65. Jonkers, I. & Lis, J. T. Getting up to speed with transcription elongation by RNA polymerase II. *Nat. Rev. Mol. Cell Biol.* **16**, 167–177 (2015).
66. Lin, C. Y. *et al.* Transcriptional amplification in tumor cells with elevated c-Myc. *Cell* **151**1215, 56–67 (2012).
67. Nie, Z. *et al.* c-Myc is a universal amplifier of expressed genes in lymphocytes and embryonic stem cells. *Cell* **151**, 68–79 (2012).
68. Littlewood, T. D., Kreuzaler, P. & Evan, G. I. All things to all people. *Cell* **151**, 11–3 (2012).
69. Gomes, S. I. L., Scott-Fordsmand, J. J. & Amorim, M. J. B. Cellular energy allocation to assess the impact of nanomaterials on soil invertebrates (enchytraeids): The effect of Cu and Ag. *Int J Environ Res Public Health* **12**, 6858–6878 (2015).
70. Scott-Fordsmand, J. J., Krogh, P. H. & Weeks, J. M. Responses of *Folsomia fimetaria* (collembola: Isotomidae) to copper under different soil copper contamination histories in relation to risk assessment. *Environ Toxicol Chem* **19**, 1297–1303 (2000).
71. OECD. *Guidelines for the Testing of chemicals No 202. Daphnia sp. Acute Immobilization Test* (Organization for Economic Cooperation and Development, Paris, 2004).
72. Chang, Y., Lye, M. L. & Zeng, H. C. Large-scale synthesis of high-quality ultralong copper nanowires. *Langmuir* **21**, 3746–3748 (2005).
73. OECD. *Guidelines for the Testing of chemicals No. 220. Enchytraeid Reproduction Test* (Organization for Economic Cooperation and Development, Paris, 2004).
74. Gomes, S. I. L., Caputo, G., Pinna, N., Scott-Fordsmand, J. J. & Amorim, M. J. B. Effect of 10 different TiO<sub>2</sub> and ZrO<sub>2</sub> (nano)materials on the soil invertebrate *Enchytraeus crypticus*. *Environ Toxicol Chem* **34**, 2409–2416 (2015).
75. R Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria (2016) URL <https://www.R-project.org/>.
76. Lemon, J. Plotrix: a package in the red light district of r. *R-News* **6**, 8–12 (2006).
77. Neuwirth, E. *RColorBrewer: ColorBrewer Palettes* <https://CRAN.R-project.org/package=RColorBrewer>. R package version 1.1-2 (2014).
78. Huber, W. *et al.* Orchestrating high-throughput genomic analysis with Bioconductor. *Nat Methods* **12**, 115–121 (2015).
79. Gautier, L., Cope, L., Bolstad, B. M. & Izrarry, R. A. Affy-analysis of Affymetrix GeneChip data at the probe level. *Bioinformatics* **20**, 307–315 (2004).
80. Durbin, J. *Distribution Theory for Tests Based on the Sample Distribution Function* (Society for Industrial and Applied Mathematics, Philadelphia, 1973).

## Acknowledgements

This work was supported by the European Union FP7 projects MODERN (Ref. 309314-2, to C.P.R., J.J.S.-F.), MARINA (Ref. 263215, to J.J.S.-F.), and SUN (Ref. 604305, to C.P.R., S.I.L.G., M.J.B.A., J.J.S.-F.), by FEDER funding through COMPETE (Programa Operacional Factores de Competitividade) and Portugal funding through FCT (Fundação para a Ciência e Tecnologia) within the project bio-CHIP (Refs FCOMP-01-0124-FEDER-041177, FCT EXPL/AAG-MAA/0180/2013, to S.I.L.G., M.J.B.A.), and by a post-doctoral grant (Ref. SFRH/BPD/95775/2013, to S.I.L.G.)

## Author Contributions

S.I.L.G., M.J.B.A. and J.J.S.-F. designed the toxicity experiment. S.I.L.G. carried out the experimental work and collected the microarray data. C.P.R. designed and implemented the novel normalization methods. C.P.R. performed the statistical analyses. All the authors jointly discussed the results. C.P.R. drafted the paper, with input from all the authors. All the authors edited the final version of the paper.

## Additional Information

**Supplementary information** accompanies this paper at <http://www.nature.com/srep>

**Competing Interests:** The authors declare no competing financial interests.

**How to cite this article:** Roca, C. P. *et al.* Variation-preserving normalization unveils blind spots in gene expression profiling. *Sci. Rep.* **7**, 42460; doi: 10.1038/srep42460 (2017).

**Publisher's note:** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2017