



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Evaluation of tools for differential gene expression analysis by RNA-seq on a 48 biological replicate experiment

Citation for published version:

Schurch, NJ, Schofield, P, Gierliski, M, Cole, C, Sherstnev, A, Singh, V, Wrobel, N, Gharbi, K, Simpson, GG, Owen-Hughes, T, Blaxter, M & Barton, GJ 2015 'Evaluation of tools for differential gene expression analysis by RNA-seq on a 48 biological replicate experiment' ArXiv.

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Evaluation of tools for differential gene expression analysis by RNA-seq on a 48 biological replicate experiment

Nicholas J. Schurch^{1,2}, Pieta Schofield^{1,2,3}, Marek Gierliński^{1,2,3}, Christian Cole^{1,2}, Alexander Sherstnev^{1,2}, Vijender Singh³, Nicola Wrobel⁶, Karim Gharbi⁶, Gordon G. Simpson^{4,7,8,9}, Tom Owen-Hughes^{3,8,10}, Mark Blaxter^{6,8,11} and Geoffrey J. Barton^{2,3,5,8,12}

¹These authors contributed equally to this work. ²Division of Computational Biology, ³Division of Gene Regulation and Expression, ⁴Division of Plant Sciences, and ⁵Division of Biological Chemistry and Drug Discovery, College of Life Sciences, University of Dundee, Dow St, Dundee, DD1 5EH, UK, ⁶Edinburgh Genomics, University of Edinburgh, Edinburgh, EH9 3JT, UK,

⁷James Hutton Institute, Invergowrie, Dundee, DD2 5DA, Scotland, UK ⁸Joint corresponding authors, ⁹g.g.simpson@dundee.ac.uk, ¹⁰t.a.owenhughes@dundee.ac.uk, ¹¹Mark.Blaxter@ed.ac.uk, ¹²g.j.barton@dundee.ac.uk

ABSTRACT

An RNA-seq experiment with 48 biological replicates in each of 2 conditions was performed to determine the number of biological replicates (n_r) required, and to identify the most effective statistical analysis tools for identifying differential gene expression (DGE). When $n_r = 3$, seven of the nine tools evaluated give true positive rates (TPR) of only 20-40%. For high fold-change genes ($|\log_2 FC| > 2$) the TPR is $> 85\%$. Two tools performed poorly; over- or under-predicting the number of differentially expressed genes. Increasing replication gives a large increase in TPR when considering all DE genes but only a small increase for high fold-change genes. Achieving a $TPR > 85\%$ across all fold-changes requires $n_r > 20$. For future RNA-seq experiments these results suggest $n_r > 6$, rising to $n_r > 12$ when identifying DGE irrespective of fold-change is important. For $n_r < 12$, superior TPR makes *edgeR* the leading tool tested. For $n_r \geq 12$, minimizing false positives is more important and *DESeq* outperforms the other tools.

RNA-seq data have been submitted to ENA archive with project ID PRJEB5348.

1 Introduction

RNA-seq has now supplanted microarrays as the technology of choice for genome-wide Differential Gene Expression (DGE) experiments. In any experimental design, selecting the appropriate number of biological replicates is a trade-off between cost and precision. For microarray methods it has been shown that low replicate experiments often have insufficient statistical power to call DGE correctly (Pan et al. 2002) and cannot accurately measure the natural biological variability (Churchill 2002). Although it is widely appreciated that increasing the number of replicates in an RNA-seq experiment usually leads to more robust results (Auer and Doerge 2010; Hansen et al. 2011; Busby et al. 2013; Liu et al. 2014), the precise relationship between replicate number and the ability to correctly identify the differentially expressed genes (i.e., the statistical power of the experiment) has not been fully explored.

The rise of RNA-seq technology has led to the development of many tools for analyzing DGE from these data (e.g., Anders and Huber 2010; Hardcastle and Kelly 2010; Robinson et al. 2010; Wang et al. 2010; Tarazona et al. 2011; Li et al. 2012; Lund et al. 2012; Trapnell et al. 2012; Li and Tibshirani 2013; Frazee et al. 2014; Law et al. 2014;

Love et al. 2014; Moulos and Hatzis 2014). Each tool makes assumptions about the statistical properties inherent to RNA-seq data and they exploit a range of normalization and analysis techniques to compute the magnitude of a DGE result and estimate its significance. Several studies have generated data specifically for the purpose of testing the assumptions intrinsic to DGE methods (Marioni et al. 2008; Consortium 2014), but most rely either on RNA-seq datasets designed to test biological hypotheses (Bullard et al. 2010; Rapaport et al. 2013; Seyednasrollah et al. 2013) or simulated data (Busby et al. 2013; Sonesson 2014), or a combination of the two (Kvam et al. 2012; Li et al. 2012; Dillies et al. 2013; Guo et al. 2013; Sonesson and Delorenzi 2013; Burden et al. 2014). The majority of studies based on analysis of experimental RNA-seq data rely on data from experiments with fewer than five replicates per condition (Marioni et al. 2008; Bullard et al. 2010; Kvam et al. 2012; Li et al. 2012; Busby et al. 2013; Dillies et al. 2013; Rapaport et al. 2013; Consortium 2014; Sonesson 2014), limiting their ability to compare the performance of DE tools as a function of replication.

Two studies explore higher replication by exploiting publically available RNA-seq data from 21 individual clones of two laboratory strains of mouse (Bottomly et al. 2011; Sonesson and Delorenzi 2013; Burden et al. 2014). Burden et al. (2014) consider false discovery rate (FDR) as the main metric for ranking five tools and conclude that at least six replicates per condition and multiplexing DGE tools gives the best results. Sonesson and Delorenzi (2013) focus on the degree of concordance between tools as a metric for comparison and conclude that none of the eleven tools they tested perform well with fewer than three replicates. Nevertheless, since the experiments are from individual mice, the data may reflect inter-individual variance in RNA expression as well as from other aspects of the experimental protocol. The same is true of studies in human that make use of data from individuals to explore higher sample replication in DGE (Guo et al. 2013; Seyednasrollah et al. 2013). Guo et al. (2013) expand replicate number by comparing six tools using RNA-seq data from breast cancer tumor-normal paired samples from fifty three individuals in The Cancer Genome Atlas (TCGA, Cancer Genome Atlas Research 2008), using this primarily to guide the construction of a simulated dataset. They conclude that all six of the tools they test suffer from oversensitivity but that *edgeR* represents the best compromise between accuracy and speed. Seyednasrollah et al. (2013) examine the performance of eight tools using mouse data (Bottomly et al. 2011) and lymphoblastoid cell data from a cohort of fifty-six unrelated Nigerian individuals from the HapMap project (International HapMap 2005). They recommend *limma* and *DESeq* for data with fewer than five replicates per condition, finding that *edgeR* is “oversensitive” and suffers from high variability in its results while *SAMSeq* suffers from a lack of statistical power with few replicates. The idea of combining DGE methods is implemented in the novel tool *PANDORA* which weights the results of different DGE tools according to their performance on test data and performs at least as well as the constituent tools (Moulos and Hatzis 2014).

In this paper the performance of DGE tools is evaluated through the first highly-replicated RNA-seq experiment designed specifically to test both the assumptions intrinsic to RNA-seq DGE tools (Gierliński et al. 2015) and to assess their performance. The paper focuses on nine popular RNA-seq specific DGE tools (as judged by citations): *baySeq*, *cuffdiff*, *DESeq*, *edgeR*, *limma*, *NOISeq*, *PoissonSeq*, *SAMSeq* and *DEGSeq* (see Table 1 for references) and assesses their performance as a function of replicate number and fold-change.

The study provides general recommendations on:

1. How many replicates future RNA-seq experiments require to maximise the sensitivity and accuracy of DGE identification and quantification.
2. The most appropriate DGE tools to use to detect DE genes in RNA-seq experiments with a given number of replicates.

Table 1: RNA-seq differential gene expression tools and statistical tests.

Name	Assumed Distribution	Normalization	Description	Version	Citations ⁴	Reference
<i>t-test</i>	normal	DEseq ¹	two-sample t-test for equal variances	-	-	-
<i>log t-test</i>	log-normal	DEseq ¹	log-ratio t-test	-	-	-
<i>Mann-Whitney</i>	none	DEseq ¹	Mann-Whitney test	-	-	Mann and Whitney (1947)
<i>Permutation</i>	none	DEseq ¹	permutation test	-	-	Efron and Tibshirani (1993)
<i>Bootstrap</i>	normal	DEseq ¹	bootstrap test	-	-	Efron and Tibshirani (1993)
<i>baySeq</i> ³	negative binomial	Internal	Empirical Bayesian estimate of posterior likelihood	1.8	109	Hardcastle and Kelly (2010)
<i>Cuffdiff</i>	negative binomial	Internal	unknown	2.1.1	481	Trapnell et al. (2012)
<i>DEGseq</i> ³	binomial	None	random sampling model using Fisher's exact test and the likelihood ratio test	1.10.0	215	Wang et al. (2010)
<i>DESeq</i> ³	negative binomial	DEseq ¹	Shrinkage variance	1.4.1	1204	Anders and Huber (2010)
<i>edgeR</i> ³	negative binomial	TMM ²	Empirical Bayes estimation & an exact test analogous to Fisher's exact test but adapted to over-dispersed data	2.2.5	822	Robinson et al. (2010)
<i>Limma</i> ³	Log-normal	TMM ²	Generalised linear model	3.4.4	15	Law et al. (2014)
<i>NOISeq</i> ³	None	RPKM	Non-parametric test based on signal-to-noise ratio	29/04/2011	113	Tarazona et al. (2011)
<i>PoissonSeq</i> ³	Poisson log-linear model	Internal	Score statistic	1.1	25	Li et al. (2012)
<i>SAMSeq</i> ³	None	Internal	Mann-Whitney test with Poisson resampling	2.0	26	Li and Tibshirani (2013)

¹ see Anders and Huber (2010), ² see Robinson and Oshlack (2010), ³ R (v2.15.1) & bioconductor (v2.10)

⁴ as reported by PubMed Central articles that reference the listed reference (28th Jan 2015)

2 Tool-specific Gold Standards

RNA was sequenced from 48 biological replicate samples of *Saccharomyces cerevisiae* in each of two well-studied experimental conditions; wild-type (WT) and a *Δsnf2* mutant. Quality control and data processing steps reject several replicates from each condition resulting in 42 WT and 44 *Δsnf2* biological replicates of 'clean' data totaling ~889M aligned reads (see Materials and Methods for a full description on the experiment, the mutant strain, the sequencing and the quality control and data processing steps).

The performance of each DGE tool as a function of replicate number and expression fold-change was evaluated by comparing the DGE results from sub-sets of these replicates against the ‘gold standard’ set of DGE results obtained for each tool with the full set of replicates. The tool-specific gold-standards were computed by running the tool on the read-count-per-gene measurements from the full set of clean data and marking as “significantly differentially expressed” (SDE) those differentially expressed genes with multiple testing corrected p -values or FDRs ≤ 0.05 . These gold-standard runs typically result in $\geq 65\%$ of the 7,126 genes in the Ensembl v68 (Flicek et al. 2011) *S. cerevisiae* annotation being identified as SDE (except for *DEGSeq* which calls considerably more genes as SDE and *NOISeq*, which calls far fewer, see Supp. Figs. S3 & S6: A).

With the tool-specific gold-standards defined, each DGE algorithm was run iteratively on i repeated sub-selections drawn from the set of clean replicates (without replacement). For each of the tools, bootstrap runs were performed with $i = 100$ iterations and $n_r = 3, \dots, 40$ replicates in each condition (*cuffdiff* was significantly slower than the other tools so the number of iterations was reduced to $i = 30$ for this tool). For a given value of n_r , the mean \log_2 transformed fold-change ($\log_2(FC)$) and median adjusted p -value or FDR calculated across all the bootstrap iterations was considered representative of the measured behavior for each individual gene. Again, genes were marked as SDE when the adjusted p -value or FDR was ≤ 0.05 . From these results, true positive, true negative, false positive and false negative rates (hereafter TPR, TNR, FPR, FNR) were then calculated as a function of n_r for four arbitrary fold-change thresholds ($|\log_2(FC)| = T \in \{0, 0.3, 1, 2\}$), by comparing the SDE genes from each bootstrap with the SDE genes from the tools gold-standard (see Materials and Methods for a detailed description of these calculations). Intrinsic to this method of measuring each tool’s performance is the assumption that the large number of replicates in the full dataset will enable each tool to unambiguously identify the ‘true’ differentially expressed genes in the experiment.

3 Results

Figure 1 shows an example of the key performance data for *edgeR* (similar figures for the other tools can be found in Supp. Figs. S2-S9). The fraction of all genes *edgeR* calls as SDE increases as a function of n_r and the impact of sampling effects on this fraction shrinks as n_r increases (Figure 1: A). The TPR performance of *edgeR* changes as a function of both

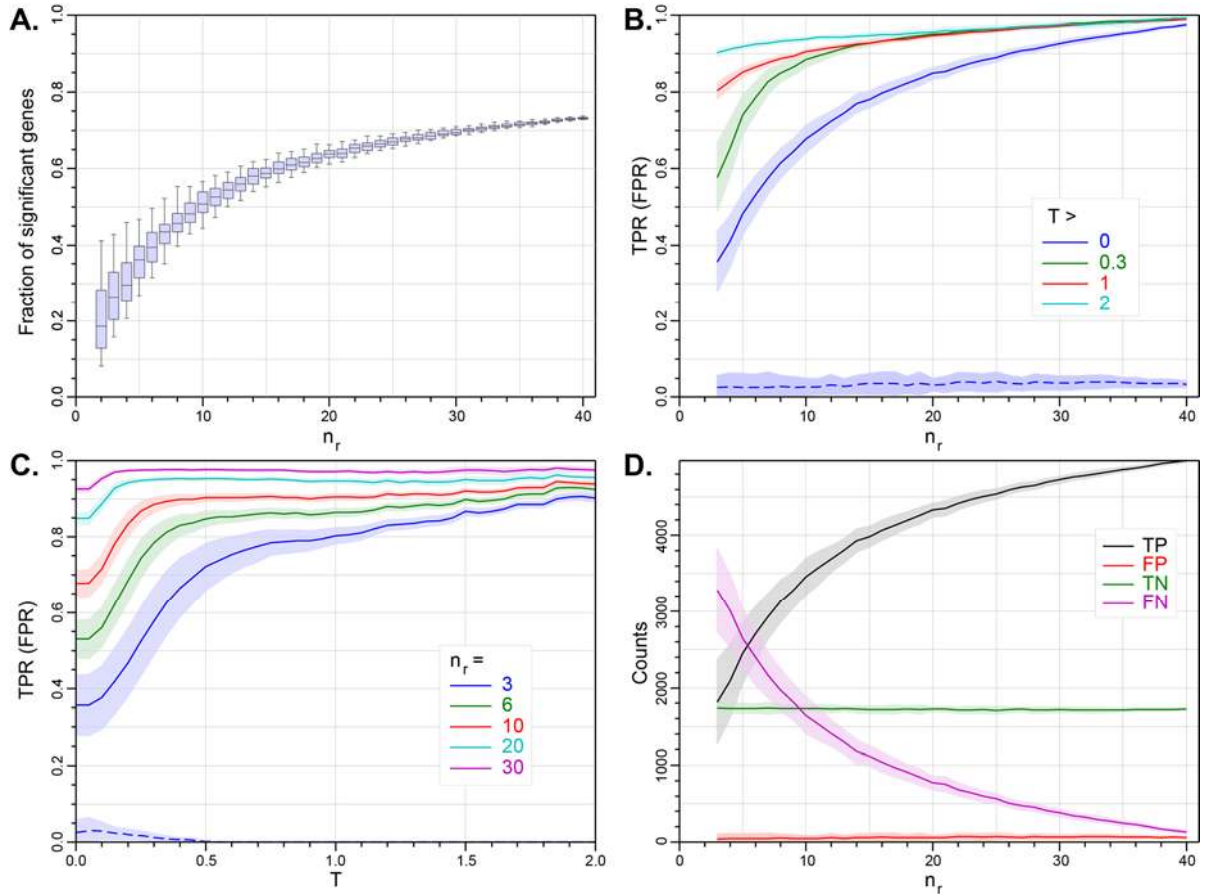


Fig 1: Statistical properties of *edgeR* as a function of $|\log_2(FC)|$ threshold, T , and the number of replicates, n_r . Individual data-points are not shown for clarity; however the points comprising the lines are each an average over 100 bootstrap iterations, with the shaded regions showing the 1-standard-deviation limits. A: The number of genes called as SDE as a function of the number of replicates (boxplots show the median, quartiles and 95% limits across replicate selections within a bootstrap run). B: mean true positive rate (TPR) as a function of n_r for four thresholds $T \in \{0, 0.3, 1, 2\}$ (solid curves, the mean false positive rate (FPR) for $T = 0$ is shown as the dashed blue curve, for comparison). Data calculated for every $\Delta n_r = 1$. C: mean TPR as a function of T for $n_r \in \{3, 6, 10, 20, 30\}$ (solid curves, again the mean FPR for $n_r = 3$ is shown as the dashed blue curve, for comparison). Data calculated every $\Delta T = 0.1$. D: The number of genes called as true/false positive/negative (TP, FP, TN and FN) as a function of n_r . The FPR remains extremely low with increasing n_r demonstrating that *edgeR* is excellent at controlling its false discovery rate. Data calculated for every $\Delta n_r = 1$.

replicate number and fold-change threshold (Figure 1: B & C). However, *edgeR* successfully controls its FDR for all combinations of both n_r and T and the primary effect of increasing the number of replicates or imposing a fold-change threshold is to increase the sensitivity of the tool, converting false negatives to true positives (Figure 1: D).

Figure 2 summarises the performance of all nine tools considered in this study as a function of replicate number and fold-change threshold. The TPR for bootstrap sub-selections with three replicates and no fold-change threshold ($n_r = 3, T = 0$) is ~ 20 -40% for all the tools except *NOISEq* and *DEGSeq*, indicating that with this few replicates these experiments were unable to identify the majority of DE genes regardless of the tool used to analyse the data (Figure 2: A).

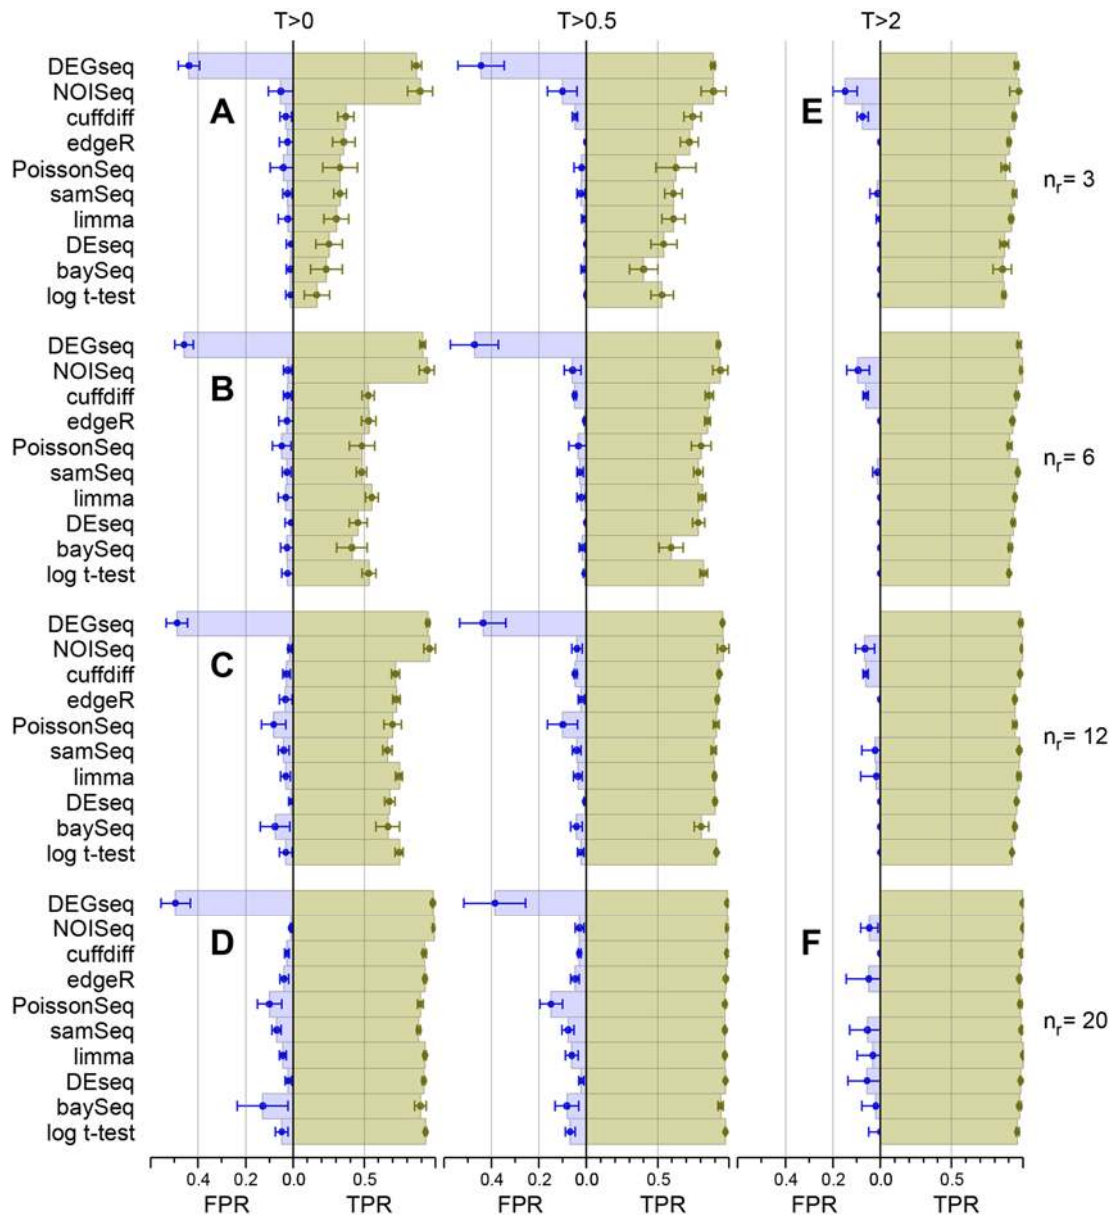


Fig 2: Comparison of the true positive rate (TPR) and false positive rate (FPR) performance for each of the DE tools on low-, medium- and highly-replicated RNA-seq data ($n_r \in \{3, 6, 12, 20\}$ – rows) for three $|\log_2(FC)|$ thresholds ($T \in \{0, 0.5, 2\}$ - columns). The TPRs & FPRs for each tool are calculated by comparing the mean number of true and false positive (TPs and FPs) calculated over 100 bootstrap iterations to the number of TPs and FPs calculated from the same tool using the full clean dataset (error-bars are 1 standard deviation). Although the TPRs and FPRs from each tool are calculated by comparing each tool against itself rather than a tool-independent ‘gold standard’ (albeit with the full clean dataset) the results are comparable across tools except for *DEGSeq* which calls a significantly larger fraction of genes as DE for all values of T and n_r – (Supp. Fig. S4) and *NOISeq* which calls a significantly smaller fraction of genes as DE for all values of T and n_r – (Supp. Fig. S7). *DEGSeq* and *PoissonSeq* produce no FPs for the highest threshold ($T > 2$) and thus no FPR is shown for them. In general, the FPR decreases with increasing fold-change threshold, and for the highest fold-change genes the FPRs are formally consistent with zero. There are exceptions however; *PoissonSeq* & *BaySeq* show an increasing FPR with increasing n_r and *cuffdiff* unexpectedly shows an increase in FPR with increasing fold-change threshold. *DESeq* appears more conservative than the other tools, consistently returning fewer FPs (particularly for high values of n_r) and fewer TPs (particularly at low values of n_r).

DEGSeq shows strong TPR performance but this is coupled with a high FPR due to *DEGSeq* overestimating the number of SDE genes regardless of the number of replicates (Supp. Fig. S4: A). *NOISeq* shows a similarly strong TPR performance with an apparently low FDR as well, however *NOISeq* identifies only a few hundred of the most strongly DE genes as SDE regardless of the number of replicates, demonstrating that this tool is insensitive to the majority of DE (Supp. Fig. S7: A). Excluding *DEGSeq* and *NOISeq*, the TPR performance for all the remaining tools is a function of fold-change (Figure 1: C & Supp. Figs. S2-S9: C). For the highest fold-change genes ($T = 2$), all the remaining seven tools show TPRs $\geq 85\%$ and (with the exception of *cuffdiff*) FPRs consistent with zero (Figure 2: E). These tools are successfully capturing the majority of the true differential expression signal for the most strongly changing genes from each tools gold standard with as few as three replicates per condition. For this cohort of high fold-change SDE genes the TPR is largely insensitive to replicate number. Irrespective of the tool, increasing the number of replicates to $n_r = 20$ for $T = 2$ provides only a modest increase in TPR from $\sim 85\%$ to $\sim 95\%$ (Figure 2: F, Figure 1: B & Supp. Figs. S2-S9: B). Increasing the number of replicates has a dramatic effect on the detection rate of genes with smaller fold-changes. Achieving an $\sim 85\%$ detection rate with *edgeR* for fold-change thresholds of $T = 1, 0.3$ & 0 requires $\sim 9, 11$ & 26 replicates, respectively (Figure 1: B & C). This behavior is reflected across all seven tools (Figure 2: A-D). Reducing the fold-change threshold reduces the TPR independently of replicate number for all the tools except *DEGSeq*, *NOISeq* and *baySeq*. The TPR performance as a function of fold-change threshold has two distinct linear regions: a shallow linear regime at high- T and a steeper region at low- T (Figure 1: C & Supp. Figs. S2-9: C). The transition between these two regions is a function of both the tool and the number of replicates. For *edgeR* with $n_r = 3$, this transition fold-change threshold is ~ 0.5 and drops to ~ 0.25 and ~ 0.15 for $n_r = 10$ & 30 respectively (Figure 1: C). These transitions represent an optimal fold-change threshold to filter the data by to maximise both the quality and the utility of the data.

The best performing tools *edgeR*, *DESeq* and *limma* successfully control their FPR, maintaining it consistently close to, or below 5% irrespective of fold-change threshold or number of replicates (Figure 1: B & C & Supp. Figs. S5 & S6: B & C) highlighting again that the primary effect of increasing replicate number is to increase the sensitivity of these tools, converting false negatives to true positives (Figure 1: D & Supp. Figs. S5 & S6: D). Other tools are not so successful in this regard but a detailed interpretation of the FPR from this test is complicated by the fact that each tool is tested against its own gold standard. A more robust method for probing the FPR performance of DE tools is presented below.

4 Tool consistency with high replicate data

The DGE tool performance tests described here assume that, given enough replicates, the tools converge on the true underlying differential expression signal in the data. This assumption was tested by clustering the DE measurements for each tool's 'gold standard' along with the results from five additional simple statistical tests applied to the same data (see Materials and Methods for a detailed description of the statistical tests). The lists of SDE genes from each tool and test (adjusted p-value or FDR threshold ≤ 0.05) were hierarchically clustered with the R package *pvclust* (Figure 3, Suzuki and Shimodaira 2006). *pvclust* uses bootstrapping to compute the statistical significance of sub-clusters within the dendrogram. Approximately Unbiased p-value percentages (AU% – Figure 3,




Fig 3: Hierarchical clustering of nine RNA-seq DE tools and five standard statistical tests using all the full clean dataset comprising 42 WT & 44 $\Delta snf2$ replicates. Lists of significantly differentially expressed (SDE) genes from each tool and test (adjusted p-value or FDR threshold ≤ 0.05) were hierarchically clustered with the R package *pvclust* (Euclidian distance metric, maximum-linkage) de. Approximately Unbiased p-value percentages (bracketed values) calculated for each branch in the clustering represent the support in the data for the observed sub-tree clustering. AU%>95% are strongly supported by the data. AU% values are not shown for branch-points where AU%=100 for clarity. The outlier clustering of *NOISeq*, *baySeq*, *DEGSeq*, *SAMSeq*, and *PoissonSeq* suggest that these tools are clearly distinct from the other tools. The 90% AU value for *DESeq* indicates that this tool clusters tightly with Cluster 1 90% of the time, and that when it does *cuffdiff* also clusters robustly with this group. Combined with the tool performance data shown in Figure 3, this suggests that, given a large number of replicates, the tools and tests in Cluster 1 (and to a lesser extent in Cluster 2) are reliably and reproducibly converging on a similar answer, and are likely to be correctly capturing the SDE signal in the data.

bracketed values) calculated for each branch in the clustering are an indication of how robust each branch is to sampling error. The most widely used tools (*edgeR* and *limma*, Table 1) are tightly grouped in a robust cluster with the standard statistical tests (Figure 3, cluster 1). In 90% of the bootstraps, *DESeq* also clusters tightly with this group, and when it does, *cuffdiff* also clusters robustly with this group, suggesting these tools and tests are reliably and reproducibly converging on approximately the same answer, given a large number of replicates (Figure 3, cluster 2). Several of the standard statistical tests

are non-parametric (Mann-Whitney, permutation and bootstrap) and use very different underlying methods compared to the tools in this cluster, indicating that the agreement of techniques within this group is not the result of a similar underlying methodology, but is likely reflective of the true differential expression signal in the data. The remaining outlier tools (*PoissonSeq*, *SAMSeq*, *NOISEq*, *DEGSeq* and *baySEQ*) cluster independently suggesting that these tools reach a different result to those in Cluster 1. Furthermore, the outlier tools cluster into three robust groups (Figure 3, clusters 3-5), indicating that they are reaching different conclusions to each other as well as to the non-outlier tools and are unlikely to be capturing the true underlying differential expression signal in the data.

5 Testing Tool False Positive Rates

Perhaps the most important performance measure for RNA-seq DE tools is their false detection rate. The large number of replicates in this study permits a simple test of the FPR for each of the tools. Two sets of n_r replicates were randomly selected (without replacement) from the WT condition. Under the null hypothesis that there is no expression change between these two sets, every gene identified as SDE is, by definition, a false positive. For each bootstrap run, the fraction of the total gene set identified as SDE was computed. The distribution of this false positive fraction as a function of the number of replicates, for each DE tool, is shown in Figure 4. This approach shows that *DEGSeq* and *SAMSeq* perform poorly even with a large number of replicates. *DEGSeq*, in particular, has poor FP performance with every bootstrap iteration identifying >5% of all genes as false positives (FPs) and a median FPR of ~50% irrespective of the number of replicates. Approximately 10% of *PoissonSeq* and *cuffdiff*, and 40% of *SAMSeq* bootstrap iterations identify >5% of all genes as FPs, suggesting that although better than *DEGSeq*, these tools are not controlling their FDR well. These results indicate that *PoissonSeq*, *cuffdiff*, *DEGSeq* and *SAMSeq* have inferior performance to the other tools considered in this study for RNA-seq DE analysis.

6 Discussion

In this work the performance of nine popular RNA-seq DE tools has been evaluated using a highly-replicated two-condition RNA-seq experiment designed specifically for the purpose of benchmarking RNA-seq DGE tools on genuine biological replicate data. Three of the nine tools, *edgeR*, *limma* and *DESeq* show excellent performance in the tests presented here. With the exception of *limma*, these are the most widely-used of the tools tested here as measured by citations (Table 1), suggesting that the majority of the RNA-seq DE analyses in the literature are using the most appropriate tools for the job. An additional important feature of all three tools is that they allow confounding experimental factors to be specified for DGE analysis. This permits *edgeR*, *limma* and *DESeq* to be used even with challenging datasets.

For experiments where it is important to capture as many of the truly SDE genes as possible but with a low number of replicates (*i.e.*, $n \lesssim 12$), the data presented here suggest *edgeR* in preference to *limma* or *DESeq* due to its superior TP identification rate. For experiments with sufficient numbers of replicates to ensure that the majority of the true SDE is already being captured (*i.e.*, $n \gtrsim 12$), and where it is, instead, important to minimise the number of false positives, the slightly better FPR performance of *DESeq*

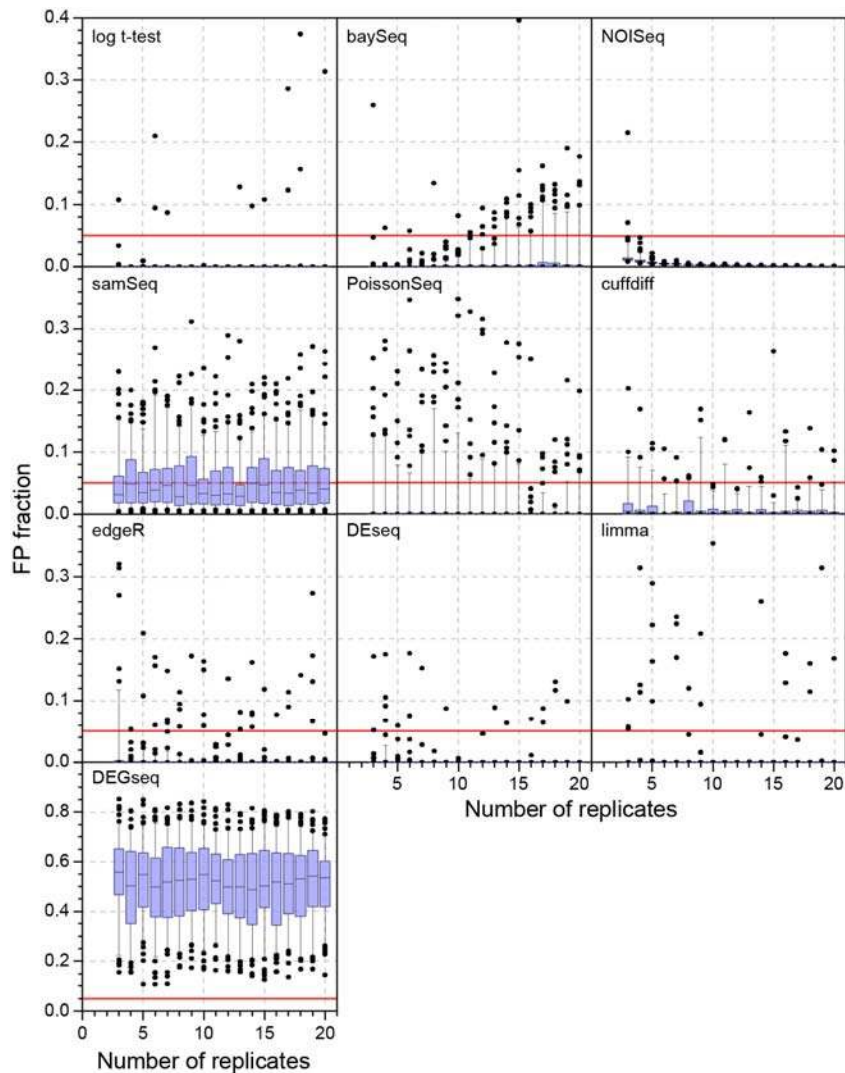


Fig 4: Testing false positive rate (FPR) performance: Each tool was used to call significantly differentially expressed (SDE) genes based on two artificial ‘conditions’ each constructed only from WT biological replicates. Genes identified as SDE are, by definition, false positives. The box plots show the median, quartiles and 95% data limits on the FPR for 100 bootstrap iterations of each of the nine tools and the log t-test for $n_r = 3, 4, \dots, 20$. The red line highlights a 5% FDR. In most cases the tools perform well for each bootstrap iteration, with only a small number of iterations showing a FPR >5%. Only *DEGSeq*, here shown with a different y-axis scale to the rest of the plots, and to a lesser extent *SAMSeq*, consistently show a high FPR, suggesting that they are struggling to control their FDR adequately.

suggests it should be the tool of choice. Conversely, *PoissonSeq*, *SAMSeq*, *NOISeq*, *DEGSeq*, *baySeq* and *cuffdiff* all show inferior performance compared to *edgeR*, *DESeq* and *limma*. Table 2 summarises recommendations for choosing RNA-seq DE tools, based on the results of these benchmarking tests. It is clear from the benchmarking runs that even the best tools have limited statistical power with few replicates in each condition, unless a stringent fold-change threshold is imposed (Figure 2). For all the tools except *baySeq*, however, the FPR, is approximately constant regardless of replicate number (Figure 4) suggesting that controlling the FNR rather than the FPR is the primary justification for imposing such a fold-change threshold. Interestingly, the variation intrinsic to the experimental procedures and protocols will result in a hard lower limit on the detectable

fold-changes for biologically relevant DE. Unfortunately, it is not possible to calculate this limit here using the gene count data alone since it requires prior knowledge of actual fold-changes to measure the impact of experimental variance.

When designing an RNA-seq experiment with the primary goal of identifying those SDE genes that change by more than a factor of two ($T = 1$), three clean replicates per condition may be sufficient. However, this is not the same as conducting the experiment with a total of three replicates, because there is a significant minority chance that one-or-more replicate within each condition should be rejected (see Gierliński et al. 2015). Conversely, for biological questions in which identifying the majority of the DE genes is important, a low-replicate experiment may not provide a sufficiently detailed view of the differential expression to inform the biology accurately. In these situations, it would be prudent to obtain at least 12 clean replicates per condition allowing the identification of $\geq 90\%$ of the truly SDE genes with $T \geq 0.3$ by any of the tools presented here except *NOISeq* or *DESeq*. It is worth recalling that identifying a gene as SDE does not necessarily equate to identifying it as biological significant and that it is important to consider both the magnitude of the measured fold-change and existing biological knowledge alongside the statistical significance when inferring a biological significance for the results of DGE experiments.

The experiment performed here is likely to be a best-case scenario and so represent the upper limit in performance of the tools tested. *S. cerevisiae* is one of the best-studied model organisms in biology, with a genome that is relatively small and well understood, few genes containing more than a single exon and there is no tissue-specific gene expression. In an organism with a more complex transcriptome, the performance of all the DE tools is likely to be worse due to the presence of multiple transcript isoforms, anti-sense non-coding RNA transcription and incomplete or poorly known annotations, particularly for 5' and 3' UTRs (Schurch et al. 2014). Although the majority of current DGE tools including the nine analysed here, rely on an existing genome annotation, the recently published DE tool *derfinder* (Frazee et al. 2014) examines DE for any region of a genome without annotations by analysing DE at base-pair resolution and grouping adjacent regions with similar DE signals. Such annotation-free differential expression tools may well represent the future for differential expression studies with RNA-seq data since they have the potential to mitigate the impact of genome annotation on detection of DE.

To the best of our knowledge the experiment presented here is the most highly replicated RNA-seq dataset to date and the only one specifically designed for testing the process of calling differential expression. As such, it will be a useful resource for the bioinformatics community as a test-bed for tool development, and for the wider biological science community as the most detailed description of transcription in wild-type & $\Delta snf2$ mutant *S. cerevisiae*. All the code for this work will be publically available (Github), and both the raw and processed data can be accessed from the European Nucleotide Archive (project ID: PRJEB5348).

Table 2: A Summary of the recommendations of this paper

	Agreement with other tools ¹	WT v WT FPR ²	Fold-change Threshold (T) ³	Tool recommended for: (# good replicates per condition) ⁴		
				<=3	<=12	>12
<i>BaySeq</i>	inconsistent	Pass				
<i>cuffdiff</i>	consistent	Fail				
<i>DEGSeq</i>	inconsistent	Fail				
<i>DESeq</i>	consistent	Pass	0			Yes
			0.5		Yes	Yes
			2	Yes	Yes	Yes
<i>edgeR</i>	consistent	Pass	0			Yes
			0.5	Yes	Yes	Yes
			2	Yes	Yes	Yes
<i>Limma</i>	consistent	Pass	0			Yes
			0.5		Yes	Yes
			2	Yes	Yes	Yes
<i>NOISeq</i>	inconsistent	Pass				
<i>PoissonSeq</i>	inconsistent	Fail				
<i>SAMSeq</i>	inconsistent	Fail				

¹ Full clean replicate dataset, see ‘Tool consistency with high replicate data’ and Figure 3.

² See ‘Testing Tool False Positive Rates’ and Figure 4

³ See ‘Differential Expression Tool Performance as a function of replicate number’

⁴ See Figure 2

7 Recommendations for RNA-seq experiment design

The results of this study suggest the following should be considered when designing an RNA-seq experiment for DGE:

- 1) At least 6 replicates per condition for all experiments.
- 2) At least 12 replicates per condition for experiments where identifying the majority of all DE genes is important.
- 3) For experiments with < 12 replicates per condition; use *edgeR*.
- 4) For experiments with > 12 replicates per condition; use *DESeq*.

- 5) Apply a fold-change threshold appropriate to the number of replicates per condition between $0.1 \leq T \leq 0.5$ (see Figure 2 and the discussion of tool performance as a function of replication).

8 Materials & Methods

8.1. $\Delta snf2$ mutant

Saccharomyces cerevisiae is one of the best studied organisms in molecular biology with a relatively small transcriptome and very limited alternative splicing and was chosen in order to give us the simplest RNA-seq data possible. *SNF2* is the catalytic subunit of ATP dependent chromatin remodeling SWI/SNF complex in yeast. *SNF2* forms part of a transcriptional activator and mutation in *SNF2* brings about significant changes in transcription (e.g., Neigeborn and Carlson 1984; Stern et al. 1984; Peterson et al. 1991; Hirschhorn et al. 1992; Peterson and Herskowitz 1992; Holstege et al. 1998; Sudarsanam et al. 2000; Becker and Horz 2002; Gkikopoulos et al. 2011; Ryan and Owen-Hughes 2011, and references therein).

8.2. *S. cerevisiae* growth conditions & RNA extraction

The *S. cerevisiae* strains used in the experiment were wildtype (BY4741 strain, WT) and $\Delta snf2$ mutant in the same genetic background. Asynchronous WT & $\Delta snf2$ mutant strains were streaked out on rich media (YPAD) to get individual colonies. For 48 replicates in both strains, single colonies were inoculated to 15 ml cultures and cells were grown to an OD600 of 0.7-0.8 (corresponding to approximately 10^6 cells) at 30°C. RNA was isolated using RNeasy mini kit (Qiagen, Manchester, UK) protocol that uses Zymolyase for yeast cell lysis and DNase treatment to remove DNA contamination. The amount of total RNA extracted ranged from 30.3 μg to 126.9 μg per sample. Although the amount RNA extracted was variable, the distributions were consistent with being drawn from the same population (Kolmogorov–Smirnov test $p = 0.16$) indicating no bias in RNA content between WT and $\Delta snf2$ mutant samples.

8.3. Library preparation, spike-in addition & sequencing

The RNA-seq experiment described here implements a ‘balanced block design’ in order to control for technical artifacts such as library batch effects (Kaisers et al. 2014), barcoding biases and lane effects via randomization of the libraries (Colbourn and Dinitz 2007; Auer and Doerge 2010). Additionally, all the replicates include artificial RNA spike-in controls in order to allow external calibration of the RNA concentrations in each sample and of the measured fold-changes between the two conditions (Jiang et al. 2011; Loven et al. 2012). The 96 samples were prepared in batches of 24 samples with 12 of each strain in each batch. Barcodes were pre-assigned randomly between the samples with barcode IDs 1-48 assigned to the $\Delta snf2$ mutant samples the 49-96 to the WT strain. For each batch the Illumina TruSeq protocol was used to prepare the sequencing library, with the addition of the ERCC spike-in standard (Ambion, Paisley, UK, Jiang et al. 2011). Briefly, samples were polyA enriched with poly-dT beads and 1 μl of 1:100 spike-in added to 19.5 μl of polyA enriched samples. Spike-in mix 1 was used with the $\Delta snf2$ mutant and

mix 2 with WT. The RNA was then fragmented and subsequently underwent both first and second strand cDNA synthesis. The cDNA then was then subjected to end repair, 3' end adenylatation, and barcode sequences were added. Finally, the un-barcoded adapters were ligated, templates purified and finally the samples were enriched via barcode-specific PCR primers. At this point the quality of the libraries was examined and passed before being diluted down to 10 nM and quantified (using fluorescence-based quantification) for accurate aliquoting for cluster generation and appropriate lane loading. Seven independent pools of the 96 barcoded samples were prepared and loaded onto seven lanes of an Illumina HiSeq 2000. Thus, each lane contains all 96 samples prepared in four batches with different spike-in mixes in each strain. The flow-cell was run for 51 cycles single-end.

8.4. Read alignment & and read-count-per-gene measurement

The lane data were de-multiplexed and processed through Cassava pipeline v1.8 to generate 672 fastq files comprising 7 technical replicates for each of the 96 biological replicates in the experiment. A total of $\sim 10^9$ reads were reported with each technical replicate having between $0.8\text{-}2.8 \times 10^6$ reads. Aggregating the technical replicates across lanes results in $\sim 10^7$ reads per biological replicate. First pass quality control of the reads was performed with *fastQC* (<http://www.bioinformatics.bbsrc.ac.uk/projects/fastqc>) for each technical replicate. The reads from each technical replicate were then aligned to the Ensembl release 68 (Flicek et al. 2011) *S. cerevisiae* genome with *bowtie2* (v2.0.0-beta7) (Trapnell and Salzberg 2009) and *TopHat2* (v2.0.5) (Trapnell et al. 2009) using the following parameters: `--max-intron-length 1000 --min-intron-length 10 --microexon-search -b2-very-sensitive --max-multihits 1`. The aligned reads were then aggregated with *htseq-count* (v0.5.3p9, Anders et al. 2014) using the Ensembl v68 *S. cerevisiae* genome annotation to give total gene read counts for all 7,126 gene features for each technical replicate. Finally, the read-count-per-gene measurements for each technical replicate were summed across sequencing lanes to give read-count-per-gene for each of the 96 biological replicates and these are then used to identify poorly correlating 'bad' replicates within the two conditions that were then subsequently removed from the analysis (see Gierliński et al. 2015 for a detailed description of this process). This resulted in a total of 42 WT & 44 $\Delta snf2$ biological replicates of 'clean' read-count-per-gene data.

8.5. Tool Details and Considerations for Differential Expression Calculations

Most of the DGE tools assessed here calculate both a fold-change (typically expressed as a logarithm to base 2 of the expression ratio, $\log_2 FC$) and a statistical significance of differential expression for each gene. The fold-change is based on the mean count across replicates in each condition and for many of the tools this includes a calculation of sample-specific normalization factors based on the gene read-count data. For this study the default normalization factors were used for each of the tools assessed. Whilst there are differences between the normalizations used by these tools it has been suggested that the details of which method is used to normalise the data does not significantly alter the downstream DGE results (Seyednasrollah et al. 2013). These normalization methods do, however, rely on the assumption that the majority of genes do not change their expression levels between conditions (e.g., Dillies et al. 2013). If this assumption is not satisfied the measurements of both DGE fold-change and significance are likely to be incorrect.

The statistical significances calculated by DGE tools are usually based on the null hypothesis of no expression change between the conditions. Calculating this significance typically relies on two key factors: 1) an assumption about the probability distribution that underlies the raw read-count measurements, and 2) being able to accurately measure the mean count and variance for each gene. Different tools assume different forms for the underlying read-count distribution including the negative binomial (*baySeq*, *Cuffdiff*, *DESeq* and *edgeR*), beta-binomial (*BBSeq*), binomial (*DEGSeq*), Poisson (*PoissonSeq*), and log-normal (*limma*) distributions. A few algorithms make no assumptions about the read-count distribution and instead take non-parametric approaches to testing for DGE (*NOISeq* & *SAMSeq*). Gierliński et al. (2015) show that for this data the majority of gene expression is consistent with both log-normal and negative binomial distributions except for the lowest expression genes, for which only the negative binomial distribution remains consistent with the data. For experiments with high numbers of replicates per condition ($n \gtrsim 12$) the mean and variance estimators can be accurately computed directly on the data. However, many RNA-seq DGE studies rely on a low numbers of replicates per condition ($n \lesssim 3$) so several of the DGE tools (e.g., *edgeR*, *limma*) compensate for the lack of replication by modelling the mean-variance relation and borrowing information across genes to shrink the given gene's variance towards the common model (Cui et al. 2005; De Hertogh et al. 2010; Robinson et al. 2010). The stabilised variance helps avoid some of the spurious false positives and negatives, but is strongly dependent on an assumed read count distribution and on the assumption that the large majority of the gene counts are not truly differentially expressed. Given these methods dependence on accurate mean and variance measurements it is somewhat surprising that scientists would contemplate doing DGE analysis without replicated data, but for completeness we note that several DGE analysis tools advertise that they can work with a single replicate per condition (Anders and Huber 2010; Robinson et al. 2010; Tarazona et al. 2011).

8.6. Bootstrap Differential Expression Calculations

A utility pipeline was written to automate the process of running each DGE algorithm iteratively on i repeated sub-selections of clean replicates. Each sub-selection is comprised of n_r replicates chosen at random without replacement (that is, an individual replicate can appear only once within each sub-selection). This bootstrapping procedure includes applying the default normalization for each tool where relevant and possible (see Tool Details and Considerations for Differential Expression Calculations) and the full output for each tool was stored in a local *sqlite* database, including the log-2 transformed fold-change and the statistical significance for every expressed gene in the *S. cerevisiae* annotation. Most of the tools return Benjamini-Hochberg (hereafter BH; Benjamini and Hochberg 1995) corrected p -values or FDRs as their measure of statistical significance. Genes with adjusted p -value or FDR ≤ 0.05 were marked as "significantly differentially expressed" (SDE). Supplementary Figure S1 shows an example of the output mean $\log_2 FC$ and median p -value data for the tool *edgeR* with $n_r = 3$.

From these data, TPRs, TNRs, FPRs and FNRs for each tool were computed as a function of the number of replicates, n_r , for four arbitrary absolute log-2 fold-change thresholds, $T \in \{0, 0.3, 1, 2\}$. A reference fold-change was used for deciding whether each gene falls above the threshold T because the measured values of mean $|\log_2 FC|$ calculated for a gene varies considerably with both the tool being used and n_r . These reference fold-changes were defined independently of the tools by applying *DESeq*

normalization (Anders and Huber 2010) to the read-count-per-gene data from the full clean set of biological replicates for each condition and then taking the log-2 transformed ratio of the mean normalized read-count-per-gene for each condition. For each individual DE calculation within a bootstrap run (i.e., an individual DE calculation with a specific tool with a given n_r), each gene was called as true/false positive/negative by comparing whether it was called as SDE in the bootstrap run, and whether it was called as SDE in the corresponding tool-specific ‘gold standard’. Then, taking each fold-change threshold in turn, the mean of number of true/false positives/negatives for genes with reference fold-changes above this threshold was calculated across all the individual DE calculations within a bootstrap run. This results in a TPR, TNR, FPR and FNR for a tool, for a given n_r and for a given T (Equations 1-4)

$$\text{TPR}(n_r, T) = \frac{\text{TP}(n_r, T)}{\text{TP}(n_r, T) + \text{FN}(n_r, T)} \quad (1)$$

$$\text{FPR}(n_r, T) = \frac{\text{FP}(n_r, T)}{\text{FP}(n_r, T) + \text{TN}(n_r, T)} \quad (2)$$

$$\text{TNR}(n_r, T) = \frac{\text{TN}(n_r, T)}{\text{TN}(n_r, T) + \text{FP}(n_r, T)} \quad (3)$$

$$\text{FNR}(n_r, T) = \frac{\text{FN}(n_r, T)}{\text{FN}(n_r, T) + \text{TP}(n_r, T)} \quad (4)$$

Uncertainties in the resulting values were calculated by propagating the standard deviations of the numbers of TPs, TNs, FPs and FNs across the calculations within each bootstrap run, to reflect the spread of calculated values due to the random sampling of replicates.

8.7. Standard statistical tests for differential expression

When assessing the performance of each DE tool on the full set of clean data we compare the tools not only within themselves, but also to the following set of standard statistical tests. For the following mathematical descriptions, $\mathbf{x}_{gk} = (x_{g1k}, x_{g2k}, \dots, x_{gn_kk})$ is a vector of n_k (clean) replicates for gene g and condition k , \bar{x}_{gk} and s_{gk}^2 are the mean and variance of this vector.

8.7.1. t-test

The null hypothesis in t-test is that the given gene under two conditions have the same mean count, $H_0: \mu_{g1} = \mu_{g2}$. We used the test statistic

$$t_g = \frac{\bar{x}_{g1} - \bar{x}_{g2}}{\sqrt{s_{g12}^2 \left(\frac{1}{n_1} + \frac{1}{n_2} \right)}} \quad (5)$$

with common variance estimator $s_{g12}^2 = [(n_1 - 1)s_{g1}^2 + (n_2 - 1)s_{g2}^2]/\nu$ and the number of degrees of freedom is $\nu = n_1 + n_2 - 2$.

8.7.2. Log-ratio t-test

This modified t-test is more appropriate for log-normally distributed data. The null hypothesis is $\ln \mu_{g1} = \ln \mu_{g2}$. The test statistic,

$$t_g = \frac{\ln \bar{x}_{g1} - \ln \bar{x}_{g2}}{\sqrt{\frac{s_{g1}^2}{n_1 \bar{x}_{g1}^2} + \frac{s_{g2}^2}{n_2 \bar{x}_{g2}^2}}} \quad (6)$$

is approximately distributed with t-distribution with $n_1 + n_2 - 2$ degrees of freedom (see Olsson 2005).

8.7.3. Mann-Whitney test

Mann-Whitney (Mann and Whitney (1947) - hereafter MW) test is a non-parametric test assessing if count rate in a gene under one condition tends to be larger than under the other. The null hypothesis is $H_0: \Pr(x_{gi1} > x_{gj2}) = \frac{1}{2}$, for each pair of replicates i and j . p-values were calculated using normal approximation (Bellera et al. 2010) and taking ties into account (Harbord (2004), p.428-431). The MW test relies on ranks, not actual data values, which makes it distribution-free. On the other hand, when every replicate in one condition is larger than every replicate in the other condition, the MW test will return the same p-value, regardless of how much the two conditions differ.

8.7.4. Permutation test

In the permutation test, counts from both conditions are pooled together (for each gene), $\mathbf{x}_g = (\mathbf{x}_{g1}, \mathbf{x}_{g2})$ and then randomly resampled B times without replacement from \mathbf{x}_g , using the original sizes, n_1 and n_2 . For the b -th random permutation $\mathbf{x}_{g1}^*(b)$ and $\mathbf{x}_{g2}^*(b)$ we find the test statistic, $D_g^*(b) = \bar{x}_{g1}^*(b) - \bar{x}_{g2}^*(b)$, which is the difference between the means of the two sampled vectors. This is compared with the observed statistic $D_g = \bar{x}_{g1} - \bar{x}_{g2}$. The test p-value is the fraction of cases where the resampled statistic exceeds the observed one, $p_g = |D_g^*(b)| > |D_g|/B$ (for more details see Efron and Tibshirani (1993), p203-219). The advantage of the permutation test is that it does not make any assumptions about the underlying distribution, but rather models it directly from data. The disadvantage is that it requires many replicates to build this underlying distribution, as it is not applicable for a typical experiment with, say, three replicates.

8.7.5. Bootstrap test

The Studentised bootstrap test described by Efron and Tibshirani (1993, p220-224) was employed here. It estimates probability distribution of the two populations under the null hypothesis of the common mean. Data are resampled with replacement to estimate the significance level. For the b -the bootstrap, $\mathbf{x}_{g1}^*(b)$ and $\mathbf{x}_{g2}^*(b)$, the test statistic is

$$t_g^*(b) = \frac{\bar{x}_{g1}^*(b) - \bar{x}_{g2}^*(b)}{\sqrt{s_{g12}^{*2}(b) \left(\frac{1}{n_1} + \frac{1}{n_2}\right)}} \quad (7)$$

where the common variance estimator is $s_{g12}^{*2}(b) = [(n_1 - 1)s_{g1}^{*2}(b) + (n_2 - 1)s_{g2}^{*2}(b)] / (n_1 + n_2 - 2)$. This is compared with the observed statistic (Equation 5). As in the permutation test, the test p-value is the fraction of cases where the resampled statistic exceeds the observed one, $p_g = |t_g^*(b)| > |t_g|/B$.

9 Acknowledgements

We acknowledge Dr Tom Walsh for his efforts and support in managing the software requirements for this project on our HPC platform and thank Dr Gabriele Schweikert for stimulating discussions. This work was supported by: The Wellcome Trust strategic awards WT09230, WT083481 & WT097945, Biotechnology and Biological Sciences Research Council (BBSRC - BB/H002286/1 & BB/J00247X/1). Additional funding for Cole, Gierliński, and Schofield - The Wellcome Trust (92530/Z/10/Z); Schurch - BBSRC grant BB/M004155/1; Sherstnev - BBSRC grant BB/H002286/1; Blaxter, Wrobel & Gharbi (GenePool) - MRC (MR/K001744/1); Owen-Hughes & Singh - The Wellcome Trust senior fellowship 095062; Prof. Simpson's lab is funded by The Scottish Government and BBSRC grants BB/M004155/1 and BB/H002286/1.

10 References

- Anders S, Huber W. 2010. Differential expression analysis for sequence count data. *Genome Biol* **11**: R106 doi:10.1186/gb-2010-11-10-r106.
- Anders S, Pyl PT, Huber W. 2014. HTSeq; A Python framework to work with high-throughput sequencing data. *bioRxiv*.
- Auer PL, Doerge RW. 2010. Statistical design and analysis of RNA sequencing data. *Genetics* **185**: 405-416 doi:10.1534/genetics.110.114983.
- Becker PB, Horz W. 2002. ATP-dependent nucleosome remodeling. *Annual review of biochemistry* **71**: 247-273 doi:10.1146/annurev.biochem.71.110601.135400.
- Bellera CA, Julien M, Hanley JA. 2010. Normal Approximations to the Distributions of the Wilcoxon Statistics: Accurate to What "N"? Graphical Insights'. *Journal of Statistics Education* **18**.
- Benjamini Y, Hochberg Y. 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society Series B (Methodological)* **57**: 289-300 doi:10.2307/2346101.
- Bottomly D, Walter NA, Hunter JE, Darakjian P, Kawane S, Buck KJ, Searles RP, Mooney M, McWeeney SK, Hitzemann R. 2011. Evaluating gene expression in C57BL/6J and DBA/2J mouse striatum using RNA-Seq and microarrays. *PloS one* **6**: e17820 doi:10.1371/journal.pone.0017820.
- Bullard JH, Purdom E, Hansen KD, Dudoit S. 2010. Evaluation of statistical methods for normalization and differential expression in mRNA-Seq experiments. *BMC bioinformatics* **11**: 94 doi:10.1186/1471-2105-11-94.
- Burden C, Qureshi S, Wilson SR. 2014. Error estimates for the analysis of differential expression from RNA-seq count data. *PeerJ PrePrints* **2**, <http://www.ncbi.nlm.nih.gov/pubmed/25337456>.
- Busby MA, Stewart C, Miller CA, Grzeda KR, Marth GT. 2013. Scotty: a web tool for designing RNA-Seq experiments to measure differential gene expression. *Bioinformatics* **29**: 656-657 doi:10.1093/bioinformatics/btt015.
- Cancer Genome Atlas Research N. 2008. Comprehensive genomic characterization defines human glioblastoma genes and core pathways. *Nature* **455**: 1061-1068 doi:10.1038/nature07385.
- Churchill GA. 2002. Fundamentals of experimental design for cDNA microarrays. *Nature genetics* **32 Suppl**: 490-495 doi:10.1038/ng1031.
- Colbourn CJ, Dinitz JH. 2007. *Handbook of combinatorial designs*. Chapman & Hall/Taylor & Francis, Boca Raton, FL.

- Consortium SM-I. 2014. A comprehensive assessment of RNA-seq accuracy, reproducibility and information content by the Sequencing Quality Control Consortium. *Nature biotechnology* **32**: 903-914 doi:10.1038/nbt.2957.
- Cui X, Hwang JT, Qiu J, Blades NJ, Churchill GA. 2005. Improved statistical tests for differential gene expression by shrinking variance components estimates. *Biostatistics* **6**: 59-75 doi:10.1093/biostatistics/kxh018.
- De Hertogh B, De Meulder B, Berger F, Pierre M, Bareke E, Gaigneaux A, Depiereux E. 2010. A benchmark for statistical microarray data analysis that preserves actual biological and technical variance. *BMC bioinformatics* **11**: 17 doi:10.1186/1471-2105-11-17.
- Dillies MA, Rau A, Aubert J, Hennequet-Antier C, Jeanmougin M, Servant N, Keime C, Marot G, Castel D, Estelle J, Guernec G, Jagla B, Jouneau L, Laloe D, Le Gall C, Schaeffer B, Le Crom S, Guedj M, Jaffrezic F. 2013. A comprehensive evaluation of normalization methods for Illumina high-throughput RNA sequencing data analysis. *Brief Bioinform* **14**: 671-683 doi:10.1093/bib/bbs046.
- Efron B, Tibshirani R. 1993. *An introduction to the bootstrap*. Chapman & Hall, New York.
- Flicek P, Amode MR, Barrell D, Beal K, Brent S, Chen Y, Clapham P, Coates G, Fairley S, Fitzgerald S, Gordon L, Hendrix M, Hourlier T, Johnson N, Kahari A, Keefe D, Keenan S, Kinsella R, Kokocinski F, Kulesha E, Larsson P, Longden I, McLaren W, Overduin B, Pritchard B, Riat HS, Rios D, Ritchie GR, Ruffier M, Schuster M, Sobral D, Spudich G, Tang YA, Trevanion S, Vandrovcova J, Vilella AJ, White S, Wilder SP, Zadissa A, Zamora J, Aken BL, Birney E, Cunningham F, Dunham I, Durbin R, Fernandez-Suarez XM, Herrero J, Hubbard TJ, Parker A, Proctor G, Vogel J, Searle SM. 2011. Ensembl 2011. *Nucleic acids research* **39**: D800-806 doi:10.1093/nar/gkq1064.
- Frazeo AC, Sabuncuyan S, Hansen KD, Irizarry RA, Leek JT. 2014. Differential expression analysis of RNA-seq data at single-base resolution. *Biostatistics* **15**: 413-426 doi:10.1093/biostatistics/kxt053.
- Gierliński M, Blaxter M, Cole C, Gharbi K, Owen-Hughes T, Schofield P, Schurch NJ, Sherstnev A, Singh V, Wrobel N, Barton GJ. 2015. Statistical models for RNA-seq data. *arXiv & submitted* doi:<http://arxiv.org/abs/1505.00588>.
- Gkikopoulos T, Schofield P, Singh V, Pinskaya M, Mellor J, Smolle M, Workman JL, Barton GJ, Owen-Hughes T. 2011. A role for Snf2-related nucleosome-spacing enzymes in genome-wide nucleosome organization. *Science* **333**: 1758-1760 doi:10.1126/science.1206097.
- Guo Y, Li CI, Ye F, Shyr Y. 2013. Evaluation of read count based RNAseq analysis methods. *BMC genomics* **14 Suppl 8**: S2 doi:10.1186/1471-2164-14-S8-S2.
- Hansen KD, Wu Z, Irizarry RA, Leek JT. 2011. Sequencing technology does not eliminate biological variability. *Nature biotechnology* **29**: 572-573 doi:10.1038/nbt.1910.
- Harbord R. 2004. Handbook of parametric and nonparametric statistical procedures, 3rd edition. *J Roy Stat Soc a Sta* **167**: 765-766 doi:DOI 10.1111/j.1467-985X.2004.331_11.x.
- Hardcastle TJ, Kelly KA. 2010. baySeq: empirical Bayesian methods for identifying differential expression in sequence count data. *BMC bioinformatics* **11**: 422 doi:10.1186/1471-2105-11-422.
- Hirschhorn JN, Brown SA, Clark CD, Winston F. 1992. Evidence that SNF2/SWI2 and SNF5 activate transcription in yeast by altering chromatin structure. *Genes & development* **6**: 2288-2298, <http://www.ncbi.nlm.nih.gov/pubmed/1459453>.

- Holstege FC, Jennings EG, Wyrick JJ, Lee TI, Hengartner CJ, Green MR, Golub TR, Lander ES, Young RA. 1998. Dissecting the regulatory circuitry of a eukaryotic genome. *Cell* **95**: 717-728, <http://www.ncbi.nlm.nih.gov/pubmed/9845373>.
- International HapMap C. 2005. A haplotype map of the human genome. *Nature* **437**: 1299-1320 doi:10.1038/nature04226.
- Jiang L, Schlesinger F, Davis CA, Zhang Y, Li R, Salit M, Gingeras TR, Oliver B. 2011. Synthetic spike-in standards for RNA-seq experiments. *Genome Res* **21**: 1543-1551 doi:10.1101/gr.121095.111.
- Kaisers W, Schwender H, Schaal H. 2014. Hierarchical clustering of DNA k-mer counts in RNA-seq fastq files reveals batch effects. *arXiv*.
- Kvam VM, Liu P, Si Y. 2012. A comparison of statistical methods for detecting differentially expressed genes from RNA-seq data. *American journal of botany* **99**: 248-256 doi:10.3732/ajb.1100340.
- Law CW, Chen Y, Shi W, Smyth GK. 2014. Voom: precision weights unlock linear model analysis tools for RNA-seq read counts. *Genome biology* **15**: R29 doi:10.1186/gb-2014-15-2-r29.
- Li J, Tibshirani R. 2013. Finding consistent patterns: a nonparametric approach for identifying differential expression in RNA-Seq data. *Statistical methods in medical research* **22**: 519-536 doi:10.1177/0962280211428386.
- Li J, Witten DM, Johnstone IM, Tibshirani R. 2012. Normalization, testing, and false discovery rate estimation for RNA-sequencing data. *Biostatistics* **13**: 523-538 doi:10.1093/biostatistics/kxr031.
- Liu Y, Zhou J, White KP. 2014. RNA-seq differential expression studies: more sequence or more replication? *Bioinformatics* **30**: 301-304 doi:10.1093/bioinformatics/btt688.
- Love MI, Huber W, Anders S. 2014. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome biology* **15**: 550 doi:10.1186/PREACCEPT-8897612761307401.
- Loven J, Orlando DA, Sigova AA, Lin CY, Rahl PB, Burge CB, Levens DL, Lee TI, Young RA. 2012. Revisiting global gene expression analysis. *Cell* **151**: 476-482 doi:10.1016/j.cell.2012.10.012.
- Lund SP, Nettleton D, McCarthy DJ, Smyth GK. 2012. Detecting differential expression in RNA-sequence data using quasi-likelihood with shrunken dispersion estimates. *Statistical applications in genetics and molecular biology* **11** doi:10.1515/1544-6115.1826.
- Mann HB, Whitney DR. 1947. On a Test of Whether One of 2 Random Variables Is Stochastically Larger Than the Other. *Ann Math Stat* **18**: 50-60 doi:DOI 10.1214/aoms/1177730491.
- Marioni JC, Mason CE, Mane SM, Stephens M, Gilad Y. 2008. RNA-seq: an assessment of technical reproducibility and comparison with gene expression arrays. *Genome Res* **18**: 1509-1517 doi:10.1101/gr.079558.108.
- Moulos P, Hatzis P. 2014. Systematic integration of RNA-Seq statistical algorithms for accurate detection of differential gene expression patterns. *Nucleic acids research* doi:10.1093/nar/gku1273.
- Neigeborn L, Carlson M. 1984. Genes affecting the regulation of SUC2 gene expression by glucose repression in *Saccharomyces cerevisiae*. *Genetics* **108**: 845-858, <http://www.ncbi.nlm.nih.gov/pubmed/6392017>.
- Olsson U. 2005. Confidence intervals for the mean of log-normal distribution. *Journal of Statistics Education* **13**.

- Pan W, Lin J, Le CT. 2002. How many replicates of arrays are required to detect gene expression changes in microarray experiments? A mixture model approach. *Genome biology* **3**: research0022, <http://www.ncbi.nlm.nih.gov/pubmed/12049663>.
- Peterson CL, Herskowitz I. 1992. Characterization of the yeast SWI1, SWI2, and SWI3 genes, which encode a global activator of transcription. *Cell* **68**: 573-583, <http://www.ncbi.nlm.nih.gov/pubmed/1339306>.
- Peterson CL, Kruger W, Herskowitz I. 1991. A functional interaction between the C-terminal domain of RNA polymerase II and the negative regulator SIN1. *Cell* **64**: 1135-1143, <http://www.ncbi.nlm.nih.gov/pubmed/2004420>.
- Rapaport F, Khanin R, Liang Y, Pirun M, Krek A, Zumbo P, Mason CE, Socci ND, Betel D. 2013. Comprehensive evaluation of differential gene expression analysis methods for RNA-seq data. *Genome biology* **14**: R95 doi:10.1186/gb-2013-14-9-r95.
- Robinson MD, McCarthy DJ, Smyth GK. 2010. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics* **26**: 139-140 doi:10.1093/bioinformatics/btp616.
- Robinson MD, Oshlack A. 2010. A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol* **11**: R25 doi:10.1186/gb-2010-11-3-r25.
- Ryan DP, Owen-Hughes T. 2011. Snf2-family proteins: chromatin remodellers for any occasion. *Current opinion in chemical biology* **15**: 649-656 doi:10.1016/j.cbpa.2011.07.022.
- Schurch NJ, Cole C, Sherstnev A, Song J, Duc C, Storey KG, McLean WH, Brown SJ, Simpson GG, Barton GJ. 2014. Improved annotation of 3' untranslated regions and complex loci by combination of strand-specific direct RNA sequencing, RNA-Seq and ESTs. *PLoS one* **9**: e94270 doi:10.1371/journal.pone.0094270.
- Seyednasrollah F, Laiho A, Elo LL. 2013. Comparison of software packages for detecting differential expression in RNA-seq studies. *Briefings in bioinformatics* doi:10.1093/bib/bbt086.
- Soneson C. 2014. compcodeR-an R package for benchmarking differential expression methods for RNA-seq data. *Bioinformatics* doi:10.1093/bioinformatics/btu324.
- Soneson C, Delorenzi M. 2013. A comparison of methods for differential expression analysis of RNA-seq data. *BMC bioinformatics* **14**: 91 doi:10.1186/1471-2105-14-91.
- Stern M, Jensen R, Herskowitz I. 1984. Five SWI genes are required for expression of the HO gene in yeast. *Journal of molecular biology* **178**: 853-868, <http://www.ncbi.nlm.nih.gov/pubmed/6436497>.
- Sudarsanam P, Iyer VR, Brown PO, Winston F. 2000. Whole-genome expression analysis of snf/swi mutants of *Saccharomyces cerevisiae*. *Proceedings of the National Academy of Sciences of the United States of America* **97**: 3364-3369 doi:10.1073/pnas.050407197.
- Suzuki R, Shimodaira H. 2006. Pvcust: an R package for assessing the uncertainty in hierarchical clustering. *Bioinformatics* **22**: 1540-1542 doi:10.1093/bioinformatics/btl117.
- Tarazona S, Garcia-Alcalde F, Dopazo J, Ferrer A, Conesa A. 2011. Differential expression in RNA-seq: a matter of depth. *Genome research* **21**: 2213-2223 doi:10.1101/gr.124321.111.
- Trapnell C, Pachter L, Salzberg SL. 2009. TopHat: discovering splice junctions with RNA-Seq. *Bioinformatics* **25**: 1105-1111 doi:10.1093/bioinformatics/btp120.

- Trapnell C, Roberts A, Goff L, Pertea G, Kim D, Kelley DR, Pimentel H, Salzberg SL, Rinn JL, Pachter L. 2012. Differential gene and transcript expression analysis of RNA-seq experiments with TopHat and Cufflinks. *Nature protocols* **7**: 562-578 doi:10.1038/nprot.2012.016.
- Trapnell C, Salzberg SL. 2009. How to map billions of short reads onto genomes. *Nat Biotechnol* **27**: 455-457 doi:10.1038/nbt0509-455.
- Wang L, Feng Z, Wang X, Wang X, Zhang X. 2010. DEGseq: an R package for identifying differentially expressed genes from RNA-seq data. *Bioinformatics* **26**: 136-138 doi:10.1093/bioinformatics/btp612.

11 Supplementary Figures




Fig S1: An example of the $\log_{10}(\text{BH-corrected p-value})$ vs $\log_2(\text{FC})$ output data from the differential expression algorithms in this study. The data shown here are from the tool *edgeR*, averaged over 100 bootstrap iterations with $n_r = 3$. *edgeR* outputs fold-changes and p-values for 6,885 (96.6%) of the 7,126 input genes, 1,571 (22.1%) of which have BH-corrected p-values ≤ 0.05 (dashed red line). Points shown in yellow represent data with values outside the axes limits.

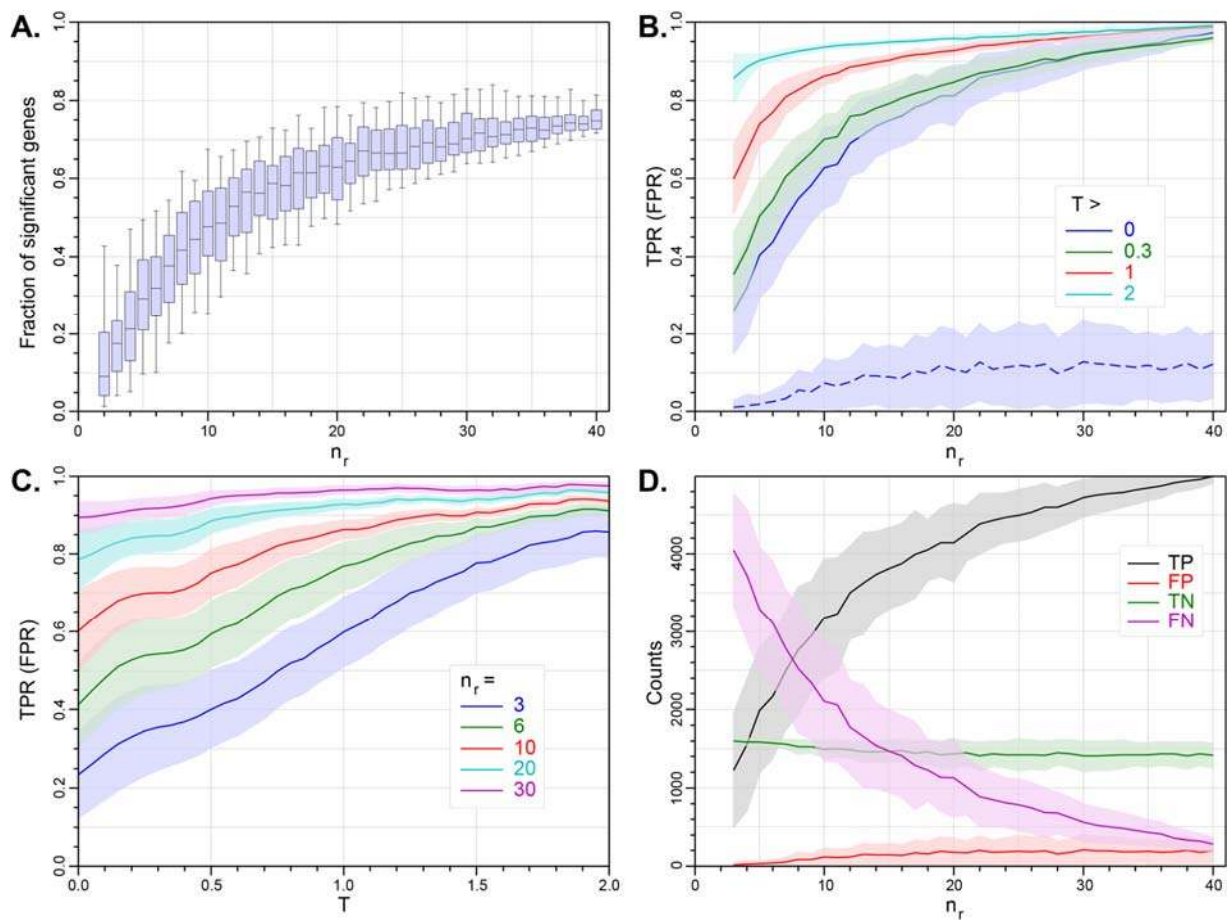


Fig S2: Statistical properties of *BaySeq* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . As in Figure 2, individual data-points are not shown for clarity; however the points comprising the lines are each an average over 100 bootstrap iterations, with the shaded regions showing the 1-standard-deviation limits. A: The number of genes called as SDE as a function of the number of replicates (boxplots show the median, quartiles and 95% data limits). B: mean TPR as a function of n_r for four fold-change thresholds $T \in \{0, 0.3, 1, 2\}$ (solid curves, the mean FPR for $T = 0$ is shown as the dashed blue curve, for comparison). Data calculated every $\Delta n_r = 1$. C: mean TPR as a function of T for $n_r \in \{3, 6, 10, 20, 30\}$ (solid curves, again the mean FPR for $n_r = 3$ is shown as the dashed blue curve, for comparison). Data calculated every $\Delta T = 0.1$. D: The number of genes called as TP, FP, TN and FN as a function of n_r .

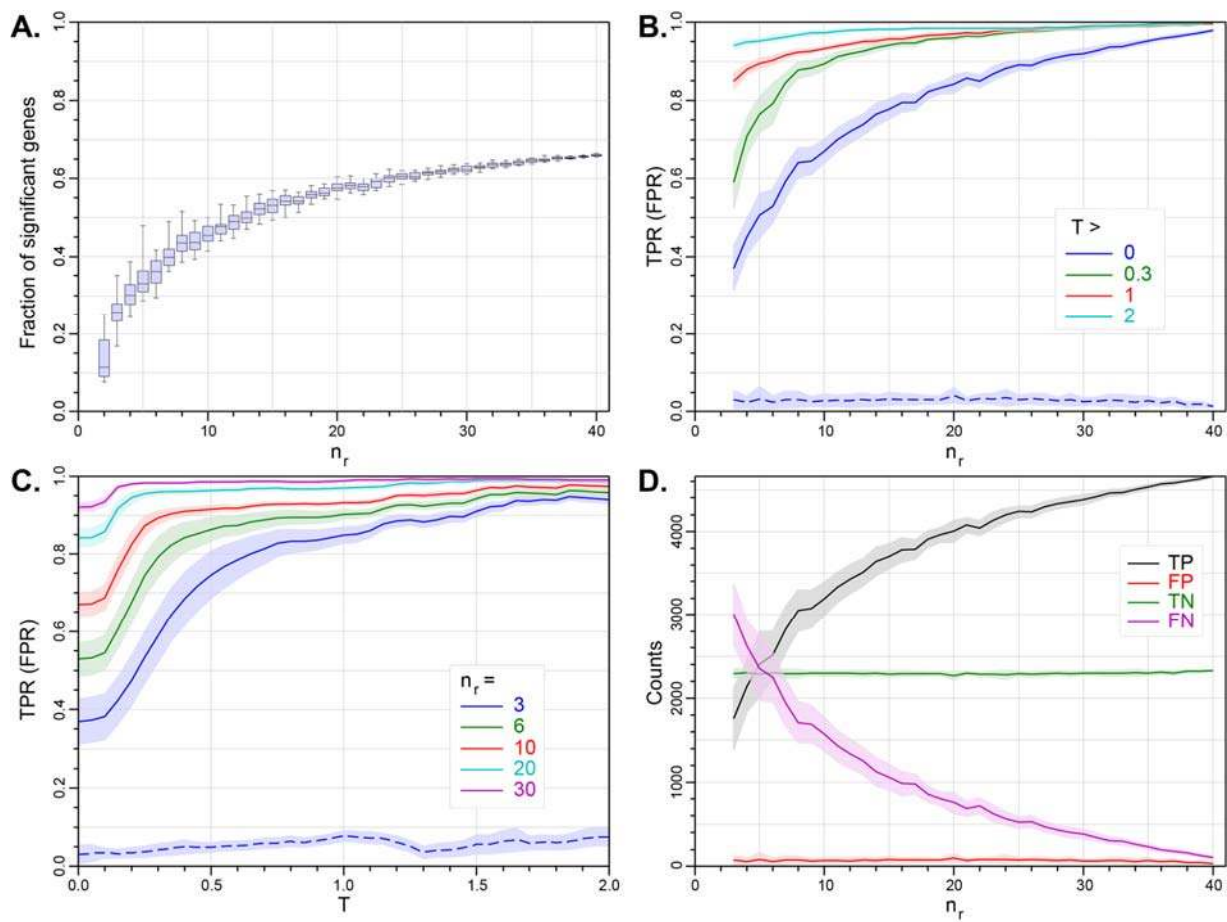


Fig S3: Statistical properties of *cuffdiff* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.

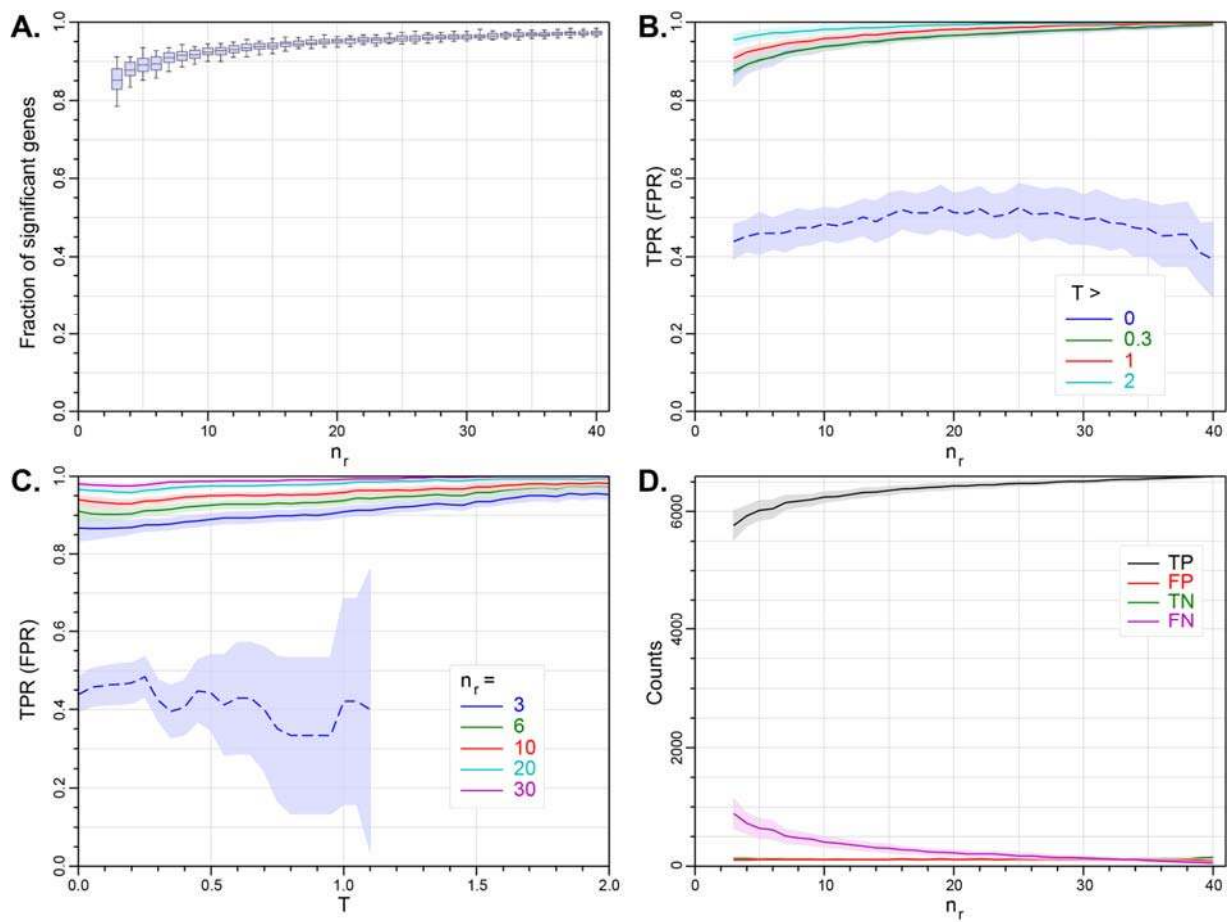


Fig S4: Statistical properties of *DEGseq* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.

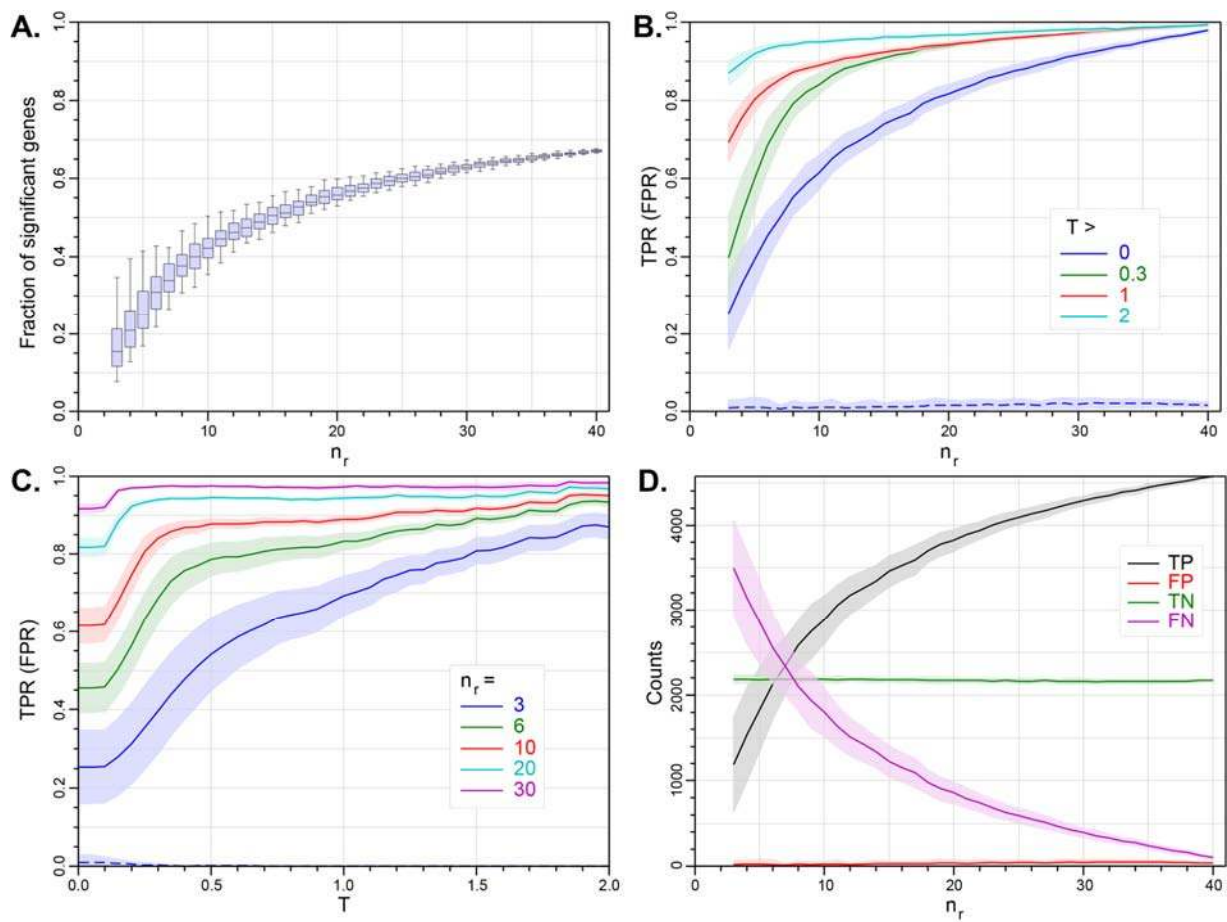


Fig S5: Statistical properties of *DEseq* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.

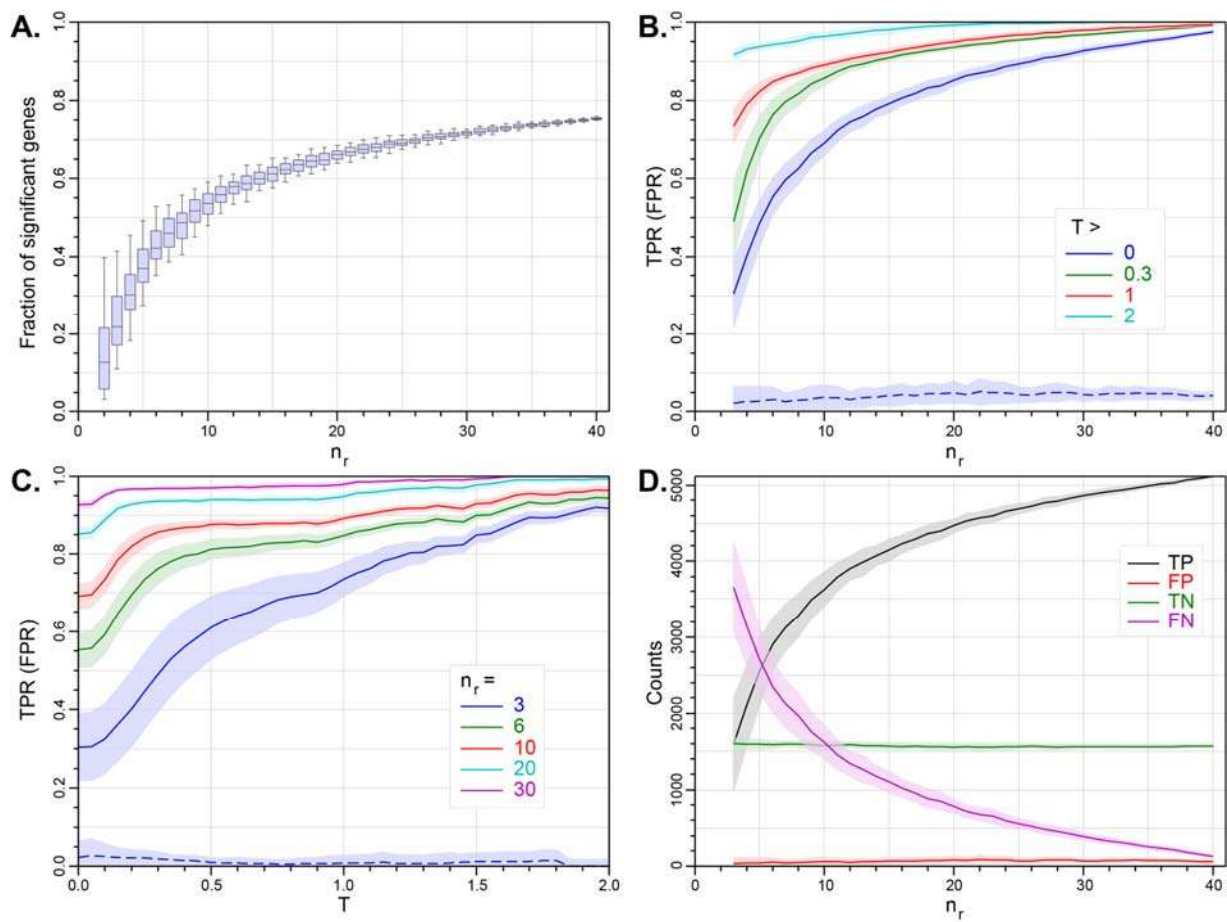


Fig S6: Statistical properties of *limma* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.

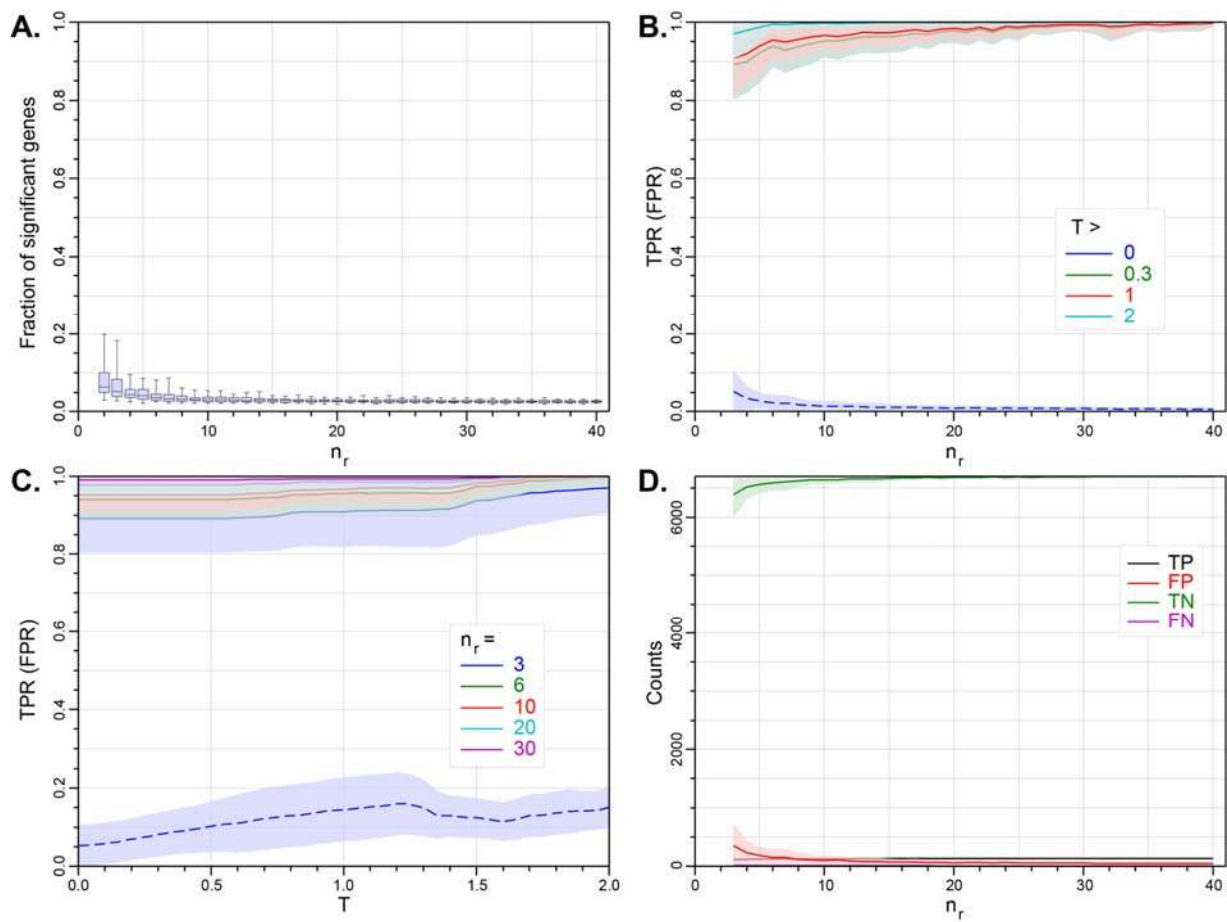


Fig S7: Statistical properties of *NOISEq* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.

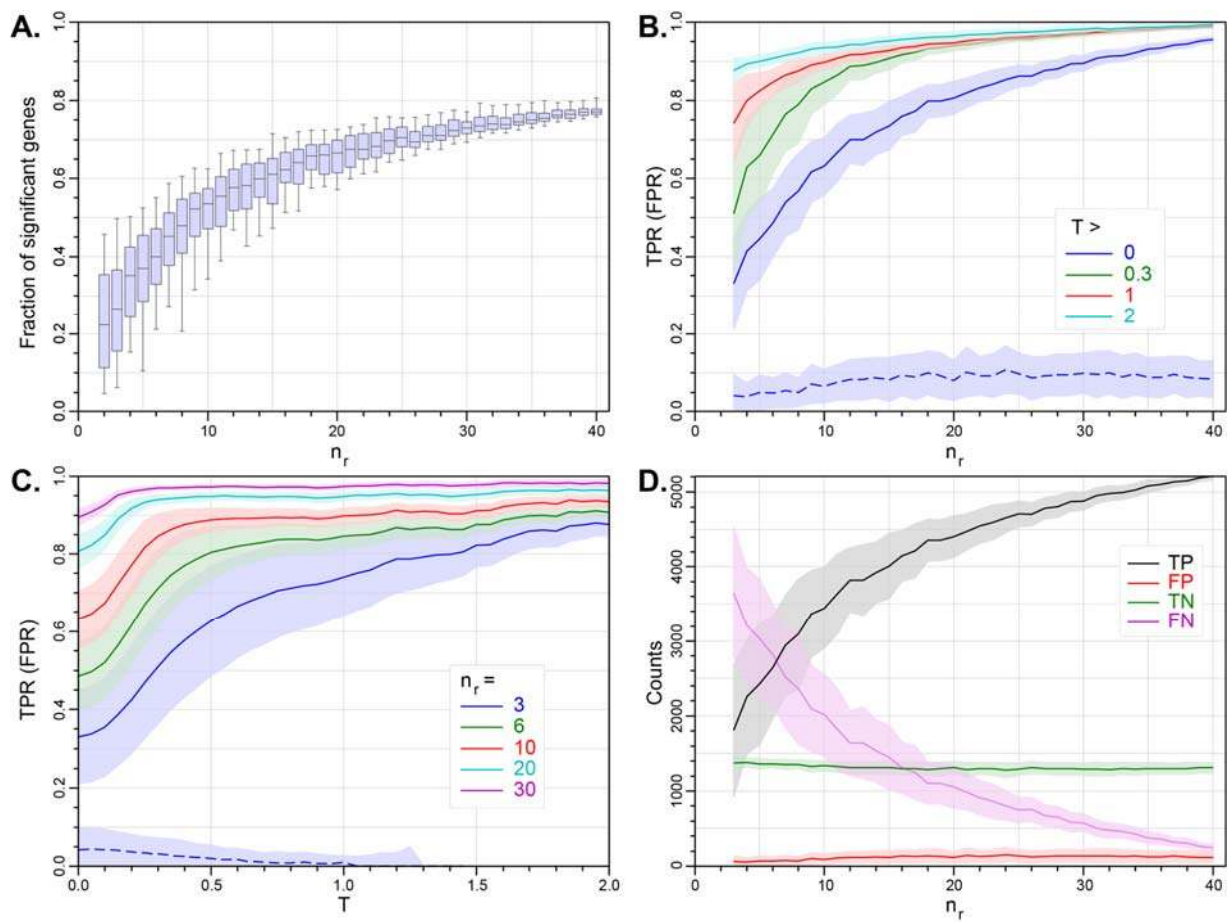


Fig S8: Statistical properties of *PoissonSeq* as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.

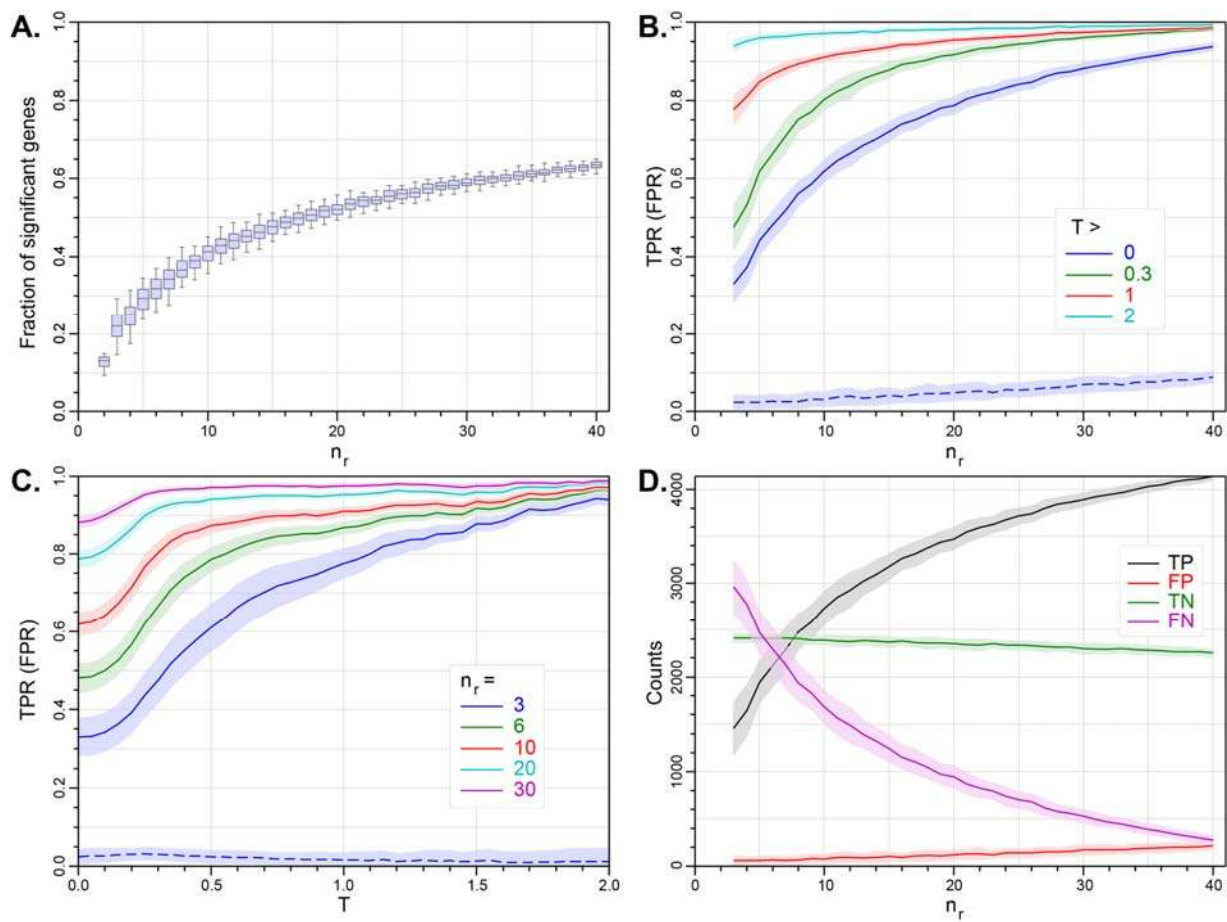


Fig S9: Statistical properties of SAMseq as a function of $|\log_2 FC|$ threshold, T and the number of replicates, n_r . See Figure S2 for detail on each panel.