Published in final edited form as:

J Bioinform Comput Biol. 2016 June ; 14(3): 1642004. doi:10.1142/S021972001642004X.

NORMALIZATION OF TRANSPOSON-MUTANT LIBRARY SEQUENCING DATASETS TO IMPROVE IDENTIFICATION OF CONDITIONALLY ESSENTIAL GENES

Michael A. DeJesus* and

Department of Computer Science, Texas A&M University, College Station, Texas, 77843, United States

Thomas R. Ioerger

Department of Computer Science, Texas A&M University, College Station, Texas, 77843, United States

Abstract

Sequencing of transposon-mutant libraries using next-generation sequencing (TnSeq) has become a popular method for determining which genes and non-coding regions are essential for growth under various conditions in bacteria. For methods that rely on quantitative comparison of counts of reads at transposon insertion sites, proper normalization of TnSeq datasets is vitally important. Real TnSeq datasets are often noisy and exhibit a significant skew that can be dominated by high counts at a small number of sites (often for non-biological reasons). If two datasets that are not appropriately normalized are compared, it might cause the artifactual appearance of differentially essential genes in a statistical test, constituting type I errors (false positives). In this paper, we propose a novel method for normalization of TnSeq datasets that corrects for the skew of read count distributions by fitting them to a Beta-Geometric distribution. We show that this read-count correction procedure reduces the number of false positives when comparing replicate datasets grown under the same conditions (for which no genuine differences in essentiality are expected). We compare these results to results obtained with other normalization procedures, and show that it results in greater reduction in the number of false positives. In addition we investigate the effects of normalization on the detection of differentially essential genes.

Keywords

Keyword1; keyword2; keyword3

1. Introduction

Sequencing of transposon-mutant libraries using next-generation sequencing has become a popular method for determining which genes and non-coding regions are essential for growth under various conditions in bacteria.¹⁹ Briefly, a transposon-mutant library is made by transfecting in a vector carrying a transposable element, such as the Himar1

*Department of Computer Science, Texas A&M University College Station, Texas, 77843, United States

transposon^{8,15}, which can insert at random locations throughout the genome (HimarI can insert randomly at any TA dinucleotide). Each mutant in the library has an insertion at a single location, but the goal is to construct a saturating library where nearly all of the potential insertion sites are represented. When grown under selective conditions, mutants with transposon insertions in essential regions will fail to survive. The abundance of the remaining insertion sites can be determined by using PCR to amplify the junctions between the transposon and the surrounding genome⁹, and the position of each insertion can be efficiently determined using a next-generation sequencer such as an Illumina HiSeq. This experiment typically yields several million reads, and the number of reads associated with each TA site is tabulated. While TA sites in non-essential regions have stochastically varying read counts, essential genes and non-coding regions (such as tRNAs, rRNAs, and sRNAs) can be identified as regions where the TA sites are uniformly devoid of insertions (i.e. read counts are 0).^{16,17,6,7}

Determining which genes in an organism are essential is a difficult problem. The primary challenge is in lower-density datasets, where the fraction of TA sites represented in the library could be in the 20–30% range. The lower the density of the dataset, the more difficult it is to determine whether a region lacks insertions due to essentiality, or just due to random statistical fluctuations. In addition, not all TA sites in an essential gene must lack insertions, as insertions can sometimes be tolerated in the N- or C-terminus of an essential gene, or in non-essential domains or linkers between domains.^{18,1} For methods that rely on comparing read-counts, the variability of the data poses an additional problem.²⁰

To address these challenges, several statistical methods have been proposed for quantifying the significance of essential genes. One method fits a Negative Binomial distribution to the insertion counts in each gene, and uses this to determine a p-value for significance of sparse regions.²¹ The length of ‘gaps’ or consecutive TA sites lacking insertions has also been used to quantify the significance of essential regions using the Extreme Value distribution.⁵ Hidden Markov Models have also been developed for analyzing TnSeq data.^{4,12} For comparison between growth conditions, the sum of read counts in a gene has been compared between conditions using a non-parametric test to identify regions with statistically significantly depressed insertions.²⁰

For methods that rely on comparison of read-counts, proper normalization of TnSeq datasets is vitally important. If two datasets that are not appropriately scaled are compared, it might cause the appearance of differentially essential genes in a statistical test, constituting type I errors (false positives). Several methods for normalizing TnSeq datasets have been proposed. Most of these methods rely on a linear transformation of the data, whereby the read-counts in a dataset are scaled by a constant factor. The simplest of these is to normalize datasets such that their read-counts have the same mean (e.g. by dividing by the total read-count). Other methods like Relative Log Expression (RLE)² and Trimmed Mean of M-values (TMM)¹⁴ have been proposed, both of which were initially developed for normalizing RNA-Seq datasets. These methods as well as others mentioned here are described in more detail in Section 2.2. Another approach is to fit a Negative Binomial distribution (or a zero-inflated Negative Binomial to help account for an abundance of empty sites) and scaling by the estimated means of the model. While scaling read-counts linearly is the most common

procedure, other methods which use a non-linear transformation have been proposed. These include quantile normalization³ which estimates empirical quantiles and then fits the datasets to match, and simulation-based normalization like the one used by ARTIST¹² which simulates a “control” dataset with similar statistical properties to an experimental dataset by sampling from a multinomial distribution.

One significant limitation of methods that linearly transform datasets is that they are susceptible to large spikes in read-counts. Because these methods multiply read-counts by a constant scalar value, they cannot reduce large outliers without also affecting small read-counts which are more common. Even if the datasets share the same mean, for instance, any skew in distribution of read-counts itself would still be present.

The distribution of read-counts in most TnSeq datasets resembles a Geometric-like distribution, in that read-counts at most sites are small (i.e. 1–50), with a (rapidly) decreasing probability of sites with large counts. Ideally, a normalization method would improve detection of conditionally essential genes between conditions by eliminating any skew and making the datasets more closely fit this Geometric-like distribution.

In this paper, we propose a novel method that corrects for the skew of read count distributions observed in many TnSeq datasets by fitting them to Geometric distribution with a variable probability parameter modeled by a Beta distribution (which we call a Beta-Geometric distribution). We show that the Beta-Geometric correction procedure (BGC) reduces the number of false positives when comparing replicate datasets grown under the same conditions (for which no genuine differences in essentiality are expected). These results are comparable to those results obtained with other normalization methods, and we show the BGC procedure produces the largest reduction in false-positives. In addition we explore the effects of BGC on the detection of differentially essential genes.

2. Beta-Geometric Correction (BGC) Normalization Method

The most common method for normalization is to divide the read counts at each TA site by the overall number of reads in a dataset, which will factor out gross differences due to the amount of data collected, analogous to the calculation of RPKMs in RNA-Seq.¹¹ A refinement of this approach that is specific to TnSeq is to scale the read counts to have the same mean over non-zero sites (which we call ‘Non-Zero Mean Normalization’ or NZMean), since different datasets can have widely varying levels of saturation, and distributing the same number of reads over fewer TA sites will naturally inflate the mean read count among them.

Despite these attempts at normalization, TnSeq datasets can still display quite different statistical patterns. In practice, some datasets appear “well behaved”, where the distribution of read counts tends to resemble a Geometric distribution (where small read-counts are most abundant, while sites with high counts are much rarer), while other datasets are skewed, with a few highly over-represented sites dominating the read-count distribution. One justification why the distribution of read-counts in (well-behaved) datasets might be expected to appear Geometrically distributed could be due to competition between the mutants in the population

of clones in the library. The abundance of the different clones in the population will vary, reflecting differences in growth rates. In the Motomura model of species abundance¹⁰, competition leads to a geometric series that describes the abundance of the species in the population, where the most fit individual has the highest abundance, and less fit individuals have exponentially decreasing abundances, with the majority of the population having very low abundance. TnSeq, by sequencing reads from this culture, is in essence obtaining a sample of read-counts in roughly the same proportion as the underlying population. Some models of abundance of populations use a Negative Binomial distribution instead. However, because the Geometric distribution is a limiting case of the Negative Binomial distribution, standardizing to a Geometric distribution can be seen as standardizing to an equivalent Negative Binomial, with size parameter $r = 1$.

The resemblance to Geometric distribution can be observed in four representative datasets shown in Figure 1. The skew away from an ideal Geometric, especially at high counts, can be seen better on a log scale (Figure 1b). These datasets are from a Himar1 Tn-mutant library in *M. tuberculosis*, where A1 and A2 are two replicates grown in vitro, and B1 and B2 representing in vivo datasets (where the library has been passaged through a mouse). Each dataset has 2 to 5 million reads distributed over 74,602 TA sites in the H37Rv genome. Datasets A1 and A2 appear to fit a Geometric distribution more closely than B1 and B2, which show greater skew. This can also be seen on a QQ-plot (quantile-quantile), where B1 and B2 veer farther away from the 1:1 diagonal than the in vitro datasets. Indeed, B1 and B2 have extremely high counts at a few individual sites (with maximum read counts of 6,009 and 16,146 respectively), compared to maximum counts of 1,693 and 1175 in the A1 and A2 datasets.

The effect of the skew observed in datasets like B1 and B2 (which is a common phenomenon in TnSeq) is that it can bias the statistical analysis of essential regions, especially for methods that depend quantitatively on the read counts. Certainly, for genes containing TA sites with high spikes in read counts, they will appear exceedingly non-essential, and it could make the gene appear differentially essential in other conditions. Simultaneously, the spikes in read counts at some TA sites will suppress the apparent level of reads at other sites, potentially making them appear relatively more essential. Figure 3, illustrates how the insertion patterns of a skewed dataset might look, before and after adjusting for the skew using the method proposed in this paper. Note that due to the non-linear nature of this transformation, high counts are significantly reduced, while sufficiently small read-counts increase.

We propose a novel method for correcting for this skew in read-count distributions by fitting each dataset to a modified distribution called a **Beta-Geometric** distribution (Equation 1), and using this to adjust the observed read counts so they more closely fit a Geometric. The Beta-Geometric distribution is like a Geometric distribution but with a variable, instead of constant, parameter p , where the variation in p is modeled by a Beta distribution. This approach is based on the observation that skewed TnSeq datasets actually appear to fit not a single Geometric with a single Bernoulli parameter, p , but the weighted sum of multiple Geometric distributions with different values of p . As weights on p , we choose the Beta distribution, with parameters ρ and κ set so that the peak is around p . The Beta distribution

has an extra degree of freedom representing dispersion around p (See Figure 4). This reflects a generative model in which individual clones in the Tn-mutant library have different growth rates, some growing slightly faster and some slightly slower than wild-type cells, depending on the location of the transposon insertion in their genome. This variability in growth rates will smear out the apparent abundance of read counts after selection (i.e. several rounds of doubling in selective conditions). In this model, the spikes in read counts would come from clones that had higher-than-average growth rates, for whatever reason (biological or random).

$$pdf(c; \rho, \kappa) = \int_0^1 Beta(p|\rho, \kappa) \times Geomtric(c|p) dp \quad (1)$$

2.1. Parameter Estimation

Given a set of read counts, Y_i at n TA sites for $i \in 1, 2, 3, \dots, n$, we assume read-counts are Geometrically distributed, with a variable parameter, p , governed by the Beta distribution:

$$\begin{aligned} Y_i &\sim Geometric(p) \\ p &\sim Beta(\kappa\rho, \kappa(1-\rho)) \end{aligned}$$

where the Beta distribution is parameterized using ρ and κ , such that ρ represents the mean of the parameter p , and κ can be thought of as analogous to a “sample size”, effectively proportional to the inverse of the variance.

We seek to estimate the parameters ρ and κ that minimize the sum of squared errors (ε) between the observed read-counts and the quantiles of the distribution:

$$\begin{aligned} \varepsilon(X; \rho, \kappa) &= \sum_i^N \left(X'_i - F^{-1}(q_i; p_i) \right)^2 \\ &= \sum_i^N \left(X'_i - \frac{\log(-q_i+1)}{\log(1-p_i)} \right)^2 \\ &= \sum_i^N \left(X'_i - \frac{\log(-q_i+1)}{\log(1-\frac{\kappa\rho-1}{\kappa-2})} \right)^2 \end{aligned} \quad (2)$$

Here, X' represents the read counts in ascending order, F^{-1} , represents the quantile function of the Geometric distribution, and $q_i \in [0, 1]$ represents the quantiles.

To facilitate the parameter estimation, the parameter ρ is estimated as $\rho = \left(\sum_i^N X_i \right)^{-1}$, which is the maximum likelihood estimate of the Geometric distribution. The remaining parameter, κ is found by determining the root of the gradient. The gradient with respect to κ is defined as follows (derivation is included in Appendix Appendix A):

$$\frac{\partial \varepsilon}{\partial \kappa} = \frac{\sum_i^N 2(2\rho-1)\log(1-q_i) \left(\log(1-q_i) - X_i \log \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right) \right)}{(\kappa-2)((\rho-1)\kappa+1)\log^3 \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right)} \quad (3)$$

The root of this gradient has an analytical solution:

$$\kappa = \frac{2 \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] - 1}{\exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] + \rho - 1}$$

Once parameters ρ and κ have been estimated, capturing the skew in the dataset, the original read counts are corrected by mapping each of them to the equivalent quantile in an ideal Geometric distribution as follows:

$$c' = F^{-1}(Q(c; \rho, \kappa); \rho) \quad (4)$$

where $Q(c; \rho, \kappa)$ is the quantile function (CDF, obtained by sampling) for the Beta-Geometric, and $F^{-1}(q; \rho)$ is the inverse of the quantile function for the Geometric distribution.

2.2. Other normalization methods

In Section 3, we compare BGC to five other normalization methods that have been proposed in the TnSeq and RNA-Seq literature.^{2,13} Because of the similarities between RNA-Seq and TnSeq procedures, as well as their dependence on normalizing count-data obtained from sequencing reads, methods used for normalizing RNA-Seq data serve as a good starting point for comparison. We include two of the most popular methods from the RNA-Seq, and as well as other methods more specific to TnSeq analysis. We briefly describe each method before presenting results.

2.2.1. Relative Log Expression (RLE)—One of the more popular normalization methods used in the RNA-Seq literature is Relative Log Expression (RLE). This normalization was proposed by Anders and Hubers and used in their DESeq method for detection of differential expression.² For each sample being normalized, RLE calculates a size-factor meant to make datasets comparable regardless of their sequence depth. The factors are calculated as follows:

$$\hat{s}_j = \text{median}_i \frac{k_{ij}}{(\prod_{v=1}^m k_{iv})^{1/m}} \quad (5)$$

where s_j represents the scaling factor for the j -th sample, and k_{ij} represents the counts at the i -th position of the j -th sample. The denominator is the geometric mean across all m

replicates, and the median over all sites (which is more robust to outliers than the mean) is taken as a scale factor for each dataset. Read-counts are then normalized by dividing them by this size-factor, rendering them comparable.

2.2.2. Trimmed Mean of M-values (TMM)—Another normalization method used in RNA-Seq is the Trimmed Mean of M-values (TMM) method. This method was developed by Robinson and Oshlack¹³, and estimates log-fold changes in expression and absolute expression:

$$M_g = \log_2 \frac{Y_{gk}/N_k}{Y_{gk'}/N_{k'}} \quad (6)$$

$$A_g = \frac{1}{2} \log_2 (Y_{gk}/N_k \times Y_{gk'}/N_{k'}) \quad (7)$$

where Y_{gk} represents counts at the g -th counts in the k -th sample, and N_k represents the total reads in that sample. The values of M_g are trimed by 30% while the samples of A_g are trimed by 5%. Finally the normalization factors are calculated by taking a weighted mean of the remaining M_g (after trimming) as follows:

$$\log_2(TMM_k^{(r)}) = \frac{\sum_{g \in G^*} w_{gk}^r M_{gk}^r}{\sum_{g \in G^*} w_{gk}^r} \quad (8)$$

where

$$M_{gk}^r = \log_2 \frac{Y_{gk}/N_k}{Y_{gr}/N_r} \quad (9)$$

and

$$w_{gk}^r = \frac{N_k - Y_{gk}}{N_k Y_{gk}} + \frac{N_r - Y_{gr}}{N_r Y_{gr}} \quad (10)$$

2.2.3. Negative Binomial—The Negative Binomial distribution (NB) is frequently used to model count data^{2,21}, particularly for data that may exhibit over-dispersion. TnSeq datasets, however, contain an overabundance of sites with read-counts of zero, representing either locations which are essential for growth or which were not sampled in the construction of the mutant library. Those libraries with a low saturation might make the mean read-count look artificially low. Ideally, the mean read-count would be calculated for

all non-essential sites, however it is difficult to separate those sites which are essential from those sites that are non-essential but missing from the library. One way to account for an excessive number of zeros, and thus attempt to separate essential sites from non-essential ones, is to use a zero-inflated model. In order to examine the influence of zeros in normalizing datasets, we compared against a zero-inflated negative binomial model (ZI-NB), which is a 2-component mixture model. The parameters were estimated by minimizing the log-likelihood of the following model:

$$P(X_i)=\pi+NB(X_i;r,p) \quad X_i=0 \quad (11)$$

$$P(X_i)=(1-\pi) \times NB(X_i;r,p) \quad X_i>0 \quad (12)$$

where π represents the probability of observing a zero outside of the negative binomial distribution, and r and p are the shape parameters of the the NB distribution. For each sample, the estimated mean of the NB distribution (i.e. $\frac{pr}{1-p}$) is used as its scaling factor.

2.2.4. Multinomial Simulation Normalization—Recently, Pritchard et al. proposed using simulation-based normalization to effectively simulate a control sample with a multinomial distribution in order to mimic the saturation (loss of library diversity) observed in the given experimental samples. This simulation method was used as part of the ARTIST pipeline for analyzing TnSeq datasets.¹²

Because this method is based on simulating samples from a multinomial distribution, it is capable of generating an arbitrary number of control samples. To compare with the other normalization methods, we took the expected value of the simulation as the normalized dataset. In addition, we simulated the dataset with the highest density to match the dataset with the lowest density. The method used in our comparison can be summarized briefly as follows:

$$\bar{C}' = E \left[Multinomial(N_x, \frac{\bar{C}}{N_c}) \right] \quad (13)$$

where $\bar{X}$ is the vector of read-counts for the input experimental sample, and $\bar{C}$ is the vector of read-counts for the input control sample, and $N_x = \sum_i X_i$ and $N_c = \sum_j C_j$, which are the total number reads in the experimental and control datasets.

2.2.5. Quantile Normalization—Another non-linear normalization method we compare against is the Quantile Normalization method (QNM). This method was proposed as a way to normalize DNA micro-array data by Bolstad et al.³ QNM normalizes datasets so that they share the same empirical distribution of values. For a given $p \times n$ matrix of counts, X_{ij} :

1. Sort each column of X , individually, to get matrix S .

2. Take the means across the rows of S and assign it to each element in the row to get S' .
3. Get the normalized matrix, X' , by rearranging each column of S' to have the same ordering as X .

This method can be seen as a special case of the transformation $x'_i = F^{-1}(G(Y_i))$, where the functions F and G are calculated empirically from the datasets being normalized.

3. Empirical Comparison of Normalization Methods

A set of 32 pairs of TnSeq datasets was obtained from various libraries of *M. tuberculosis* Tn-mutants grown under different conditions, with each condition being tested in duplicate. The raw read counts were reduced to unique template counts using sequencing barcodes⁹, though we will continue to refer to them generically as ‘read counts’ throughout this paper. Each dataset had an average of 2.4M total counts, with a range of 1.1–5.4M. Densities (i.e. fraction of TA sites represented in each dataset) were in the range of 38% to 69%.

The Beta-Geometric correction (BGC) was applied to each of the 64 datasets (followed by NZMean normalization). As an example, Table 1 contains statistics for the original datasets A1, A2, B1 and B2 (corresponding to the ‘in vitro’ and ‘Trans02c’ datasets among the 32 pairs), as well as the values of ρ and κ estimated by the BGC method. The dispersion parameter κ is lower for the B1 and B2 datasets, consistent with the greater variability that is observed in those datasets. A QQ-plot of the corrected values for dataset B2 is shown in Figure 5, displaying a much better fit to the Geometric distribution, with the skew removed (compare to Figure 2).

One empirical metric we can use to evaluate whether our correction method helps is to compare replicate datasets. In two datasets selected from the same Tn-mutant library under the same growth conditions, one would ideally expect no differences in essentiality of genes. However, in practice, there is usually high variability observed in TnSeq datasets, even between biological replicates. Any method for statistical analysis of TnSeq data has to be conservative enough not to detect many differentially essential genes between replicates. Yet, when using a permutation test (described below) on multiple pairs of replicates, we often observe differentially essential genes, in some cases far beyond what would be expected from random statistical sampling differences. We attribute many of these false positives to the skew inherent in individual datasets. Our goal in this paper is to show that, by fitting each dataset to a Beta-Geometric distribution, we can correct for the skew in read counts, and thereby reduce many of these false positives. This enhanced normalization method could be applied to other TnSeq analysis methods to improve the detection of statistically significant differentially essential genes.

3.1. Permutation Test to Identify Conditionally Essential Genes

In order to evaluate the differential essentiality of a gene between two conditions, possibly with multiple replicates of each, we use a non-parametric permutation test on the corrected and normalized counts at TA sites within the gene. Briefly, the counts are summed over all sites in a gene and replicates to determine the mean in each condition and then the difference

is compared to background distribution of means from 10,000 random permutations of the sites. The p-value is calculated from the number of times the observed mean is greater than one of the samples.

Suppose we have m_1 replicates (datasets) in condition A, and m_2 replicates in condition B. Let C_{ij} be a $(m_1 + m_2) \times n$ matrix of counts at each of n TA sites i within the gene, for each dataset j .

$$\Delta = \frac{1}{n|A|} \sum_{j \in A} \sum_i^n C_{ij} - \frac{1}{n|B|} \sum_{j \in B} \sum_i^n C_{ij} \quad (14)$$

Ten thousand random permutations of the counts in matrix C_{ij} are generated, and Δ' is calculated for each permutation. The p-value is estimated as the number of times $\Delta > \Delta'$ (or $\Delta < \Delta'$ for negative differences).

3.2. Reduction in Type I Errors

To assess the impact of the different normalization procedures when performing a comparative analysis of TnSeq datasets, we compared replicate datasets against each other. Because the datasets in each pair of replicates are selected under the same condition, the expectation is that there should be no differentially essential genes between them. A *false positive* was defined as a gene that had a $p < 0.05$, since no statistically significant differences in essentiality are expected between replicates of the same growth condition. Note that because of the large number of genes in the *M. tuberculosis* genome (i.e. 3,989), the permutation test is expected to incorrectly reject the null hypothesis on as many as 5% of the genes through chance alone.

Table 3 presents the number of false positives obtained by using the permutation test after normalizing with the different methods. Using NZMean normalization as a reference, an average of 71.4 false positives are detected over the 32 pairs of datasets. BGC reduces false-positives in 22 out of 32 cases. In comparison to other methods, BGC reduces the most false positives in 14.8 out of 32 conditions (fraction due to ties), which is more than any other normalization method. The next best normalization method was RLE, achieving the greatest reduction of false positives in 7.7 datasets. On average, BGC reduces the number of false-positives the most, achieving a mean reduction of 21.7 Type I errors overall.

No method achieves a consistent reduction in the number of false-positives on all datasets. However, even though false-positives are increased in some datasets, the amount of false-positives increased by BGC is generally small (i.e. average of 6.4). In addition, most normalization methods tend to increase false-discoveries on the same conditions, suggesting these conditions are problematic for most of the methods. For instance, on condition Trans01c, which was the condition that proved toughest for BGC (increasing false positives by 23), most other methods increased false positives as well. RLE increased false positives by 11, and TMM by 141. Only ZI-NB reduced false positives by two.

Because of the way BGC corrects for the skew in datasets, it is most likely to have a more substantial effect on those cases where there is a large skew between datasets. Table 2 contains some statistics for the datasets for which applying BGC resulted in the largest reduction in read-counts (BXD04), and in vitro (where the false positives were nearly unchanged). As can be seen, the condition on which BGC performed the best showed a very high skew and kurtosis (3rd and 4th moments of read-counts) between its replicates, where as the skew and kurtosis in the in vitro datasets were much smaller by comparison. For comparison, the skew of a dataset fitting an ideal Geometric distribution would approximately 2.0 (depends slightly on the mean). The skew in the in vitro datasets is quite close to this value, implying they are not very skewed. By correcting the skew in datasets and adjusting them to a Geometric distribution (with a variable parameter), the BGC will have more success in those datasets that are more highly skewed. On those datasets where the read-counts are not skewed, BGC is expected to have less of an effect, but these are likely the datasets that would benefit the least from normalization (as is the case for the in vitro datasets).

3.3. Effect on Detection of Differential Essentiality

So far, the previous sections have focused on the effects of BGC on reducing the number of false-positives when comparing replicates of the same condition (where no true positives are expected). It is important, however, to study the effects of BGC on detecting genes when the datasets are grown on different conditions (and thus at least some differentially essential genes, or true positives, are expected). Determining the effects of normalization on detecting true positives is complicated by the fact that it is difficult to determine a (complete) set of genes which are known a priori to be differentially essential in the conditions studied. This renders a proper analysis of the true-positive rate between normalization methods prohibitively difficult.

Instead, to study the effects of the normalization method on the comparative analysis between conditions, each pair of replicates for all the *in-vivo* conditions was compared against the pair of replicates grown *in vitro*. This way we can get an idea of how the normalization methods would affect the overall number of significant hits (though we cannot say for certain whether this leads to more true positives or not). Table 4 contains the total number of genes labeled as differentially essential (relative to in-vitro) after normalizing with each of the procedures. Differentially essential (DE) genes were those which were assigned an adjusted p-value of $q < 0.05$ (using the Benjamini-Hochberg correction for multiple comparisons). On average, the TMM method tended to predict more genes as differentially essential, with a mean of 406 DE genes, followed by RLE with a mean of 398. On the other hand, MSN showed a tendency to consistently predict the least number of DE genes, predicting an average of 67 genes as DE. The BGC method falls in between, predicting an average of 253 genes as DE.

To further explore the effect of normalizing with the BGC method, we plotted the number of DE genes detected before and after applying BGC normalization (See Figure 6). A slight reduction in the number of DE genes identified is seen in most conditions (possibly representing a decrease in the number of false-positives obtained by correcting for the skew).

This shows that the reduction in false positives between replicates is not achieved at the cost of a dramatic reduction in overall DE genes detected between conditions. However, when the number of genes classified as DE is low (due to possible under-detection of true positives), the BGC procedure tends to increase the number of DE genes predicted. On the other hand, when the number of DE genes predicted is exceedingly high (> 500), BGC normalization significantly decreases the number of DE genes predicted. This phenomenon suggests that applying BGC adjusts datasets so that they produce results that are less extreme in terms of number of DE genes detected.

4. Discussion

Analysis of TnSeq data has become a valuable tool for determining differentially essential genes. However, the large amount of intrinsic variability that is observed in these datasets (e.g. read-counts) makes direct comparison between datasets problematic. Common ways of normalizing the datasets have focused primarily on a linear transformation of read counts between datasets^{2,14}, usually by making their mean-read counts comparable. While important, normalization of the means alone is not enough to correct for the large skew that is observed in some datasets.

Other non-linear normalization methods have been proposed in the past to overcome the limitations of scaling datasets by a constant factor.^{3,12} Indeed, the BGC method is similar to quantile normalization³, except traditional quantile normalization scales datasets together based on an empirical distribution function, without making assumptions about the form of the distribution. On the other hand, the simulation-based approach taken by ARTIST is fundamentally different.¹² It attempts to simulate the effects of selection on the control dataset, by sampling read-counts from a multinomial distribution to obtain a new, simulated, control sample that has approximately the same number of reads and saturation.

We proposed the BGC method for adjusting datasets for comparative analysis. This method showed the largest overall reduction in false-positives out of all the normalization methods studied. What sets BGC apart from most of the other methods evaluated is the fact that it is a non-linear transformation of the data that is based on adjusting observed reads to an ideal distribution. It assumes that the skew in read counts comes from dispersion in the parameter p underlying a Geometric distribution. The skew is captured by fitting the data to a Beta-Geometric distribution, which allows the parameter p of the Geometric distribution to vary according to a Beta distribution. The original read counts are then adjusted back to an ideal Geometric distribution by matching quantiles. This approach is nonlinear, with high-counts (spikes) being reduced and unusually suppressed counts increased. We choose to correct read counts back to a Geometric distribution (with a variable parameter), since such a profile of abundances at different TA sites (i.e. high proportion of low counts, low proportion of high counts) would be expected from sampling from a population of competing cells with a range of growth rates.

In addition to reducing false positives in replicate datasets from the same condition, we examined the effects of applying BGC when comparing datasets of different conditions (where at least some true positives are expected). While it is difficult to say with certainty

how the BGC method affects the detection of true differentially essential genes, we showed that in most cases it tends to decrease the number of differentially essential genes slightly, likely due to reducing false-positives. As the overall reduction was relatively small, this suggests that the reduction of Type I errors that is seen when comparing replicates of the same condition is not does not come at the expense of a large reduction in the overall number of positives detected.

One potential limitation of BGC, along with most of the normalization methods examined here (except ZI-NB and MSN), is that they do not take the saturation (or density) of the data into account when adjusting reads. Accounting for different saturation levels is especially important when comparing datasets from different libraries, where saturation levels can be significantly imbalanced due to differences in biological selection. ZI-NB and MSN take into consideration the differences in saturation of the libraries in their own ways (ZI-NB by using a mixture model to allow the Negative Binomial distribution to include some, but not all, empty sites; and MSN by adjusting the saturation of the control dataset). Despite this limitation, BGC actually produces a larger reduction in false-positives compared to ZI-NB and MSN. This suggests that correcting for the skew in datasets may be more important for reducing false-positives than accounting for the difference in saturation, particularly for the well-saturated datasets such as those examined here (with insertion densities in the range of 38% to 69%). A future direction for this work could be to modify BGC so that it takes into consideration the differences in saturation levels between datasets.

Acknowledgments

This work was supported by NIH grant U19 AI107774 (TRI). We thank Chris Sassetti for providing the TnSeq datasets used in this study.

References

1. Akerley BJ, Rubin EJ, Camilli A, Lampe DJ, Robertson HM, Mekalanos JJ. Systematic identification of essential genes by in vitro mariner mutagenesis. *Proc Natl Acad Sci USA*. Jul.1998 95:8927–8932. [PubMed: 9671781]
2. Anders S, Huber W. Differential expression analysis for sequence count data. *Genome Biol*. 2010; 11(10):R106. [PubMed: 20979621]
3. Bolstad BM, Irizarry RA, Astrand M, Speed TP. A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics*. Jan; 2003 19(2):185–193. [PubMed: 12538238]
4. DeJesus MA, Ioerger TR. A Hidden Markov Model for identifying essential and growth-defect regions in bacterial genomes from transposon insertion sequencing data. *BMC Bioinformatics*. 2013; 14:303. [PubMed: 24103077]
5. DeJesus MA, Zhang YJ, Sassetti CM, Rubin EJ, Sacchettini JC, Ioerger TR. Bayesian analysis of gene essentiality based on sequencing of transposon insertion libraries. *Bioinformatics*. Mar; 2013 29(6):695–703. [PubMed: 23361328]
6. Gawronski, Jeffrey D., Wong, Sandy MS., Giannoukos, Georgia, Ward, Doyle V., Akerley, Brian J. Tracking insertion mutants within libraries by deep sequencing and a genome-wide screen for haemophilus genes required in the lung. *PNAS*. 2009; 106(38):16422–16427. [PubMed: 19805314]
7. Griffin, Jennifer E., Gawronski, Jeffrey D., DeJesus, Michael A., Ioerger, Thomas R., Akerley, Brian J., Sassetti, Christopher M. High-resolution phenotypic profiling defines genes essential for mycobacterial growth and cholesterol catabolism. *PLoS Pathog*. Sep.2011 7(9):e1002251. [PubMed: 21980284]

8. Lampe DJ, Churchill ME, Robertson HM. A purified mariner transposase is sufficient to mediate transposition in vitro. *The European Molecular Biology Organization Journal*. 1996; 15(19): 5470–5479.
9. Long, JE., DeJesus, M., Ward, D., Baker, RE., Ioerger, TR., Sassetti, CM. Identifying essential genes in *Mycobacterium tuberculosis* by global phenotypic profiling. In: Lu, Long Jason, editor. *Methods in Molecular Biology: Gene Essentiality*. Vol. 1279. Springer; 2015.
10. McGill BJ, Etienne RS, Gray JS, Alonso D, Anderson MJ, Benecha HK, Dornelas M, Enquist BJ, Green JL, He F, Hurlbert AH, Magurran AE, Marquet PA, Maurer BA, Ostling A, Soykan CU, Ugland KI, White EP. Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. *Ecol Lett*. Oct; 2007 10(10):995–1015. [PubMed: 17845298]
11. Mortazavi A, Williams BA, McCue K, Schaeffer L, Wold B. Mapping and quantifying mammalian transcriptomes by RNA-Seq. *Nat Methods*. Jul; 2008 5(7):621–628. [PubMed: 18516045]
12. Pritchard JR, Chao MC, Abel S, Davis BM, Baranowski C, Zhang YJ, Rubin EJ, Waldor MK. ARTIST: high-resolution genome-wide assessment of fitness using transposon-insertion sequencing. *PLoS Genet*. Nov.2014 10(11):e1004782. [PubMed: 25375795]
13. Robinson MD, McCarthy DJ, Smyth GK. edgeR: a Bioconductor package for differential expression analysis of digital gene expression data. *Bioinformatics*. Jan; 2010 26(1):139–140. [PubMed: 19910308]
14. Robinson MD, Oshlack A. A scaling normalization method for differential expression analysis of RNA-seq data. *Genome Biol*. 2010; 11(3):R25. [PubMed: 20196867]
15. Rubin, Eric J., Akerley, Brian J., Novik, Veronica N., Lampe, David J., Husson, Robert N., Mekalanos, John J. In vivo transposition of mariner-based elements in enteric bacteria and mycobacteria. *PNAS*. 1999; 96(4):1645–1650. [PubMed: 9990078]
16. Sassetti, Christopher M., Boyd, Dana H., Rubin, Eric J. Genes required for mycobacterial growth defined by high density mutagenesis. *Molecular Microbiology*. 2003; 48(1):77–84. [PubMed: 12657046]
17. Sassetti, Christopher M., Rubin, Eric J. Genetic requirements for mycobacterial survival during infection. *PNAS*. 2003; 100(22):12989–12994. [PubMed: 14569030]
18. Smith V, Chou KN, Lashkari D, Botstein D, Brown PO. Functional analysis of the genes of yeast chromosome V by genetic footprinting. *Science*. Dec.1996 274:2069–2074. [PubMed: 8953036]
19. van Opijnen T, Camilli A. Transposon insertion sequencing: a new tool for systems-level analysis of microorganisms. *Nat Rev Microbiol*. Jul; 2013 11(7):435–442. [PubMed: 23712350]
20. Zhang YJ, Ioerger TR, Huttenhower C, Long JE, Sassetti CM, Sacchettini JC, Rubin EJ. Global assessment of genomic regions required for growth in *Mycobacterium tuberculosis*. *PLoS Pathog*. Sep.2012 8(9):e1002946. [PubMed: 23028335]
21. Zomer A, Burghout P, Bootsma HJ, Hermans PW, van Hijum SA. ESSENTIALS: software for rapid analysis of high throughput transposon insertion sequencing data. *PLoS ONE*. 2012; 7(8):e43012. [PubMed: 22900082]

Biographies

Michael A. DeJesus received his bachelors degree in computer science in 2008 from the University of Puerto Rico, Mayaguez. He received a masters degree in computer science from Texas A&M University in 2012. Currently, he is a Ph.D. in Computer Science at Texas A&M. His research focuses on using machine learning and statistical pattern recognition techniques to analyze sequence data.

Thomas R. Ioerger graduated with honors from The Pennsylvania State University in 1989, securing a B.S. in Molecular and Cell Biology. He received an M.S. and Ph.D. in Computer Science from the University of Illinois in Urbana-Champaign, the latter in 1996. He is an associate professor in the Department of Computer Science and Engineering at Texas A&M

University. His primary research interests are in the areas of bioinformatics and machine learning..

Appendix A. Derivation

To minimize the SSE, we find the root of the derivative of SSE with respect to κ :

$$\frac{\partial SSE}{\partial \kappa} = \frac{\sum_i^N 2(2\rho-1)\log(1-q_i) \left(\log(1-q_i) - X_i \log \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right) \right)}{(\kappa-2)((\rho-1)\kappa+1)\log^3 \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right)} = 0$$

To facilitate finding the root we ignore the denominator and remove constant terms from the numerator as these do not affect the final result:

$$\begin{aligned} & \sum_i^N \log(1-q_i) \left(\log(1-q_i) - X_i \log \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right) \right) = 0 \\ & \sum_i^N \log^2(1-q_i) = \sum_i^N X_i \log(1-q_i) \log \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right) \\ & \frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} = \log \left(\frac{-\rho\kappa+\kappa-1}{\kappa-2} \right) \\ & \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] = \frac{-\rho\kappa+\kappa-1}{\kappa-2} \\ & (-\kappa+2) \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] = \rho\kappa-\kappa+1 \\ & -\kappa \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] + 2 \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] = \rho\kappa-\kappa+1 \\ & 2 \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] - 1 = \kappa \left(\rho - 1 + \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] \right) \\ & 2 \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] - 1 \\ & \kappa = \frac{2 \times \exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] - 1}{\exp \left[\frac{\sum_i^N \log^2(1-q_i)}{\sum_i^N X_i \log(1-q_i)} \right] + \rho - 1} \end{aligned}$$

Fig. 1.

a) Histogram of non-zero read counts obtained from *M. tuberculosis* Tn-mutant libraries. A1, A2 are replicates grown in vitro, and B1 and B2 are replicates grown in vivo. The black line represents a Geometric fit. b) Histogram of read counts on a log scale.

Fig. 2.
QQ-plot of the raw read counts for dataset B2, and the theoretical Geometric quantiles.

Fig. 3.

Example insertion pattern in a before and after adjusting spikes in read-counts. Unusually large read-counts can cause regions to appear to be differentially essential, artificially deflating counts at other sites below the mean (dashed line). Using a non-linear transformation, large spikes are decreased while low counts are increased, adjusting them to be more in line with each other.

a)

b)

Fig. 4.

a) Example of a Beta distribution with $\rho = 0.05$, and $\kappa = 40$. b) Histogram of counts from a regular Geometric distribution ($p = 0.05$, black curve), and a Beta-Geometric distribution ($\rho = 0.05$, $\kappa = 40$, red).

Fig. 5. QQ-plot of the raw read counts for dataset B2, and the Beta-Geometric variables obtained by sampling the parameter p from a Beta distribution with estimated parameters $\rho = 0.0078$ and $\kappa = 434.7$.

Fig. 6.

Scatter plot of the number of the differentially essential (DE) obtained with and without applying BGC. The solid black line represents the identity line. Applying BGC results in a reduction in the number of DE genes identified in most conditions, possibly representing a reduction in false-positives. In addition, BGC produces results which are less extreme, increasing the number of DE genes identified when this number is low and decreasing the number of DE genes identified when it is very high.

Table 1

Fitting of parameters for example datasets.

Data-set	Total Reads	Insertion Density	Mean Count	Max Count	ρ	κ
A1	3.12M	49.3%	84.7	1,693	0.0118	911.1
A2	1.93M	52.6%	49.2	1,175	0.0203	493.9
B1	2.78M	41.1%	89.8	6,009	0.0111	422.0
B2	3.65M	38.1%	128.4	16,146	0.0078	434.7

Table 2

Effect of skew on the change in the number of false positives (relative to NZMean) after applying BGC, for three representative conditions.

Dataset	Density	NZMean	Skew	Kurtosis	False Positives
BXD04 replicate 1	43.8	51.5	44.8	3997.8	-287
BXD04 replicate 2	54.0	86.2	7.9	183.8	
in vitro replicate 1	49.3	84.7	3.2	19.6	2
in vitro replicate 2	52.6	49.2	2.9	16.6	
Trans01c replicate 1	58.3	37.8	7.3	164.4	23
Trans01c replicate 2	65.6	51.4	5.3	61.9	

Table 3

Change in the Number of Type I Errors relative to the NZMean method. False positives are defined as genes with $p < 0.05$ under the permutation test between replicates of the same condition. Methods marked with † have been normalized with the Non-Zero Mean (NZMean) method after performing the corresponding normalization. The normalization methods compared were Beta-Geometric Correction (BGC), Relative Log Expression (RLE), Trimmed Mean of M-Values (TMM), Zero-Inflated Negative Binomial (ZI-NB), multinomial simulation normalization (MSN), and Quantile Normalization (QNM). Values which show the largest reduction in false positives for each condition are in bold. Mean reduction and the number of times each method achieves the best correction are shown at the bottom (ties are weighted by total number of methods with sharing score)

Condition (pair of replicates)	NZMean	RLE	TMM	ZI-NB	MSN†	QNM†	BGC†
AJ	13	-2	1	2	0	1	-3
BL6	74	-20	-21	-12	1	-12	-25
BXD01	2	0	1	-1	1	1	0
BXD03	0	0	0	0	0	0	2
BXD04	535	-328	364	-142	2	-73	-287
BXD05	33	1	99	1	0	-1	13
BXD06	91	-4	82	-6	0	8	-10
BXD07	78	-17	-21	-7	2	-10	-42
BXD08	241	-75	-105	-42	-9	-52	-154
BXD09	6	0	11	0	0	0	-3
CAST	17	-3	-2	1	2	0	-3
CCcont	2	0	98	0	0	0	3
DS01	12	-2	14	0	-1	-2	-4
DS02	22	2	13	-1	1	-1	-1
DS04	49	3	79	2	3	2	-17
DS0c	42	-9	-1	-5	-1	-2	-15
GP01	74	-44	-41	-1	-4	-17	-46
<i>in vitro</i>	2	0	77	0	0	0	2
PWK	100	3	4	0	2	-3	-3
Trans01	32	-13	-1	-5	-6	-2	-13
Trans01c	62	11	141	-2	1	2	23
Trans02c	84	-37	-33	-11	-1	-12	-28
Trans03	46	1	101	-4	-4	-3	3

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Condition (pair of replicates)	NZMean	RLE	TMM	ZI-NB	MSN [†]	QNM [†]	BGC [†]
Trans03c	52	0	121	2	1	2	5
Trans05	142	-7	1	-5	-1	2	-19
Trans05c	30	3	64	4	1	2	12
Trans07	158	-11	43	-10	-1	-1	-61
Trans07c	70	-8	-4	-2	1	0	1
Trans09	32	2	27	-4	0	0	-1
Trans09c	78	-27	341	7	1	-1	-3
Trans11	22	-9	3	-6	1	0	-4
Trans11c	85	21	446	-5	4	-2	-18
Mean	71.4	-17.8	59.4	-7.9	-0.13	-5.4	-21.7
Num. of Best	N/A	7.7	0.2	4.5	2.2	2.5	14.8

Table 4

Number of genes classified as differentially essential (DE) by the permutation test after applying the different normalization methods. DE genes are defined as genes with $q < 0.05$ under the permutation test between a pair replicates of the given condition and a pair replicates grown in-vitro. Methods marked with † have been normalized with the Non-Zero Mean (NZMean) method after performing the corresponding normalization. The normalization methods compared were Beta-Geometric Correction (BGC), Relative Log Expression (RLE), Trimmed Mean of M-Values (TMM), Quantile Normalization (QNM), Zero-Inflated Negative Binomial (ZI-NB), and multinomial simulation normalization (MSN)

Condition (vs in vitro)	# DE GENES					
	NZMean	RLE	TMM	ZI-NB	MSN	BGC
AJ	441	486	652	436	37	436
BL6	383	323	249	303	282	367
BXD01	366	372	421	355	37	347
BXD03	330	432	355	321	11	330
BXD04	315	273	266	302	13	307
BXD05	308	381	315	301	38	306
BXD06	356	338	412	346	36	337
BXD07	301	388	315	298	38	299
BXD08	299	416	258	288	37	289
BXD09	329	535	320	338	43	326
CAST	460	478	819	461	36	461
DS01	387	553	400	381	43	384
DS02	379	654	371	367	37	382
DS04	336	334	511	315	35	334
DS0c	323	431	326	324	37	329
GP01	844	481	852	628	545	763
PWK	453	478	559	428	33	436
Trans01	307	268	423	286	43	308
Trans01c	36	61	44	38	37	35
Trans02c	398	290	253	306	284	380
Trans03	266	257	425	255	37	259
Trans03c	91	84	223	77	35	87
Trans05	283	841	351	277	35	274

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Condition (vs in vitro)	# DE GENES						
	NZMean	RLE	TMM	ZI-NB	MSN	QNM	BGC
Trans05c	149	1208	833	86	37	152	272
Trans07	282	734	409	272	39	277	226
Trans07c	39	142	35	42	42	45	100
Trans09	524	278	477	446	3	497	137
Trans09c	24	34	72	27	40	25	74
Trans11	695	307	1164	563	2	535	154
Trans11c	43	76	86	33	34	35	83
Mean Num. of DE	325	398	407	297	67	312	253