Assessment of pharmacologic area under the curve when baselines are variable

Authors: Jeremy D. Scheff [1], Richard R. Almon [2,3], Debra C. DuBois [2,3], William J. Jusko [3], and Ioannis P. Androulakis [1,4,*]

Affiliations: [1] Biomedical Engineering Department, Rutgers University, Piscataway, NJ, 08854, USA; [2] Department of Biological Sciences, State University of New York at Buffalo, Buffalo, NY, 14260, USA; [3] Department of Pharmaceutical Sciences, State University of New York at Buffalo, Buffalo, NY, 14260, USA; and [4] Chemical and Biochemical Engineering Department, Rutgers University, Piscataway, NJ, 08854, USA

* Corresponding author:
Ioannis P. Androulakis
E-mail: yannis@rci.rutgers.edu

Abstract

Purpose: The area under the curve (AUC) is commonly used to assess the extent of exposure of a drug. The same concept can be applied to generally assess pharmacodynamic responses and the deviation of a signal from its baseline value. When the initial condition for the response of interest is not zero, there is uncertainty in the true value of the baseline measurement. This necessitates the consideration of the AUC relative to baseline to account for this inherent uncertainty and variability in baseline measurements.

Methods: An algorithm to calculate the AUC with respect to a variable baseline is developed by comparing the AUC of the response curve with the AUC of the baseline while taking into account uncertainty in both measurements. Furthermore, positive and negative components of AUC (above and below baseline) are calculated separately to allow for the identification of biphasic responses.

Results: This algorithm is applied to gene expression data to illustrate its ability to capture transcriptional responses to a drug that deviate from baseline and to synthetic data to quantitatively test its performance.

Conclusions: The variable nature of the baseline is an important aspect to consider when calculating the AUC.

Keywords: AUC, Baseline, Pharmacogenomics, Microarrays, Bioinformatics

Introduction

In pharmacology, the area under the plot of plasma concentration of a drug versus time after dosage (called "area under the curve" or AUC) gives insight into the extent of exposure to a drug and its clearance rate from the body. By integrating over time rather than looking at individual concentration measurements, a more accurate estimate of the overall exposure to the drug is obtained (Kwon 2001). Such measurements have also been found to be meaningful for assessing the net pharmacologic response to a given dose of drug (Krzyzanski and Jusko 1998).

More generally, a similar measure can be calculated for any quantitative response that deviates from its baseline measurement. For instance, when analyzing gene expression data, one might not be interested only in which genes are activated or suppressed; it could be more informative to also consider the extent to which gene expression is perturbed for full periods of time. Many existing algorithms for analyzing gene expression data do not take the time scale into account, similar to an ANOVA where each time point is treated as a different condition (Tusher, Tibshirani et al. 2001); others take the ordering of points into account, but not the length of time between points (Ernst and Bar-Joseph 2006). This could be problematic because many experiments have high sampling frequencies immediately after the experimental perturbation (*e.g.* drug administration) and much lower sampling frequencies as the response measure slowly recovers. Using the AUC to evaluate the extent of the transcriptional response of a gene does account for the temporal ordering of samples and the period of time between samples.

When the AUC is calculated for an exogenously administered drug that is not endogenously produced, it is known that the initial concentration is zero and, eventually, the drug will be eliminated and the concentration will return to zero. However, if there is some nonzero baseline value for the response of interest, it becomes less clear how to accurately calculate the AUC. For instance, gene expression values are not typically zero under normal conditions. But even if this is taken into account, the baseline cannot be expressed by a single constant value. It also has some variability. In particular, a gene regulated by the circadian clock has widely different expression values throughout the day, but even in a case where a gene expression profile is flat, normal biological noise will perturb the gene transcript abundance from its mean value. Because pharmacologic experiments are typically performed with limited numbers of replicates, incorporating variability in baseline values into the AUC calculation becomes important.

Thus, we have developed an algorithm to calculate the AUC for a pharmacologic effect such as gene expression relative to a variable baseline estimate, which can be used to discover significant net responses such as transcriptional regulatory effects in gene expression data. This allows for the segregation of responses into categories representing up-regulated, down-regulated, and biphasic responses that take into account both positive and negative changes in values. Its performance is assessed by running simulations on synthetic data. Additionally, the method is applied to real gene expression data studying the response to acute corticosteroid treatment in rat liver (Jin, Almon et al. 2003) to test its ability to select biologically relevant genes.

Materials and Methods

The general idea of the proposed method is to estimate both the overall AUC (called "AUC") and the baseline AUC (called "baseline") for each pharmacologic measurement, and then to compare these values to determine if the AUC significantly deviates from the baseline. Thus, first we define how these two values are calculated.

Estimating the baseline and its error

A simple, ideal experiment would produce frequently-sampled and precise data from both a treated group and control group, facilitating the direct comparison between these two groups. Real experiments are often constrained by limited resources so that they do not meet this ideal standard; for instance, a control group may only be available at a single initial time point rather than at every experimental time point. Depending on the system being studied and the available experimental data, and assuming no measurements are made at negative time points, there are generally three different situations in which the baseline may be estimated using different methods:

1. Baseline is estimated from measurements at only t=0
2. Baseline is estimated from measurements at t=0 and $t=T_{last}$, for values that return to baseline by the end of the experiment
3. Baseline is estimated from a separate control group with measurements collected at each time point

For the first definition, if no separate control group is available and the only baseline measurement is taken before treatment, baseline can only be estimated from the values at this single time point. Then, the baseline is computed by assuming the response stays constant at its original mean value from the first time point through the last time point. Then, the area under this flat line gives the baseline, and its variability can be expressed by using the standard deviation of those measurements. This makes sense for data that do not generally return to baseline, such as during chronic dosing of a drug in which late measurements are typically different than early measurements.

The second definition applies in the case where data exhibits a perturbation followed by a return to baseline, such as in acute dosing of a drug when enough time is given to ensure that washout occurs. Measurements from the beginning and the end of the time series can be used to estimate the true baseline. In these cases, the baseline is estimated by averaging the replicates at the first and last time points and finding the area under the line between them. The error in the baseline is calculated from the standard deviation of these replicates. The main advantage of this relative to the first definition is that experiments typically have a low number of replicates for each measurement, so by taking the first and last time points into account a significantly more accurate estimate for the baseline can be obtained.

The third definition can only be applied in cases where measurements for a separate control group are available at each time point. If these measurements are available, then they represent an excellent estimate of the baseline. In the case where the baseline varies (for example, when regulated by circadian rhythms), the ratio of the treated and control AUCs can be computed (Wolfsegger 2007).

Estimating the AUC and its error

The AUCs and confidence intervals are calculated by bootstrapping (Bonate 1998). This method is described in the following steps:

1. For each data point, sample with replacement from all of the data points measuring the same response (*i.e.* gene expression value, drug concentration, etc.) at the same time.
2. Calculate the AUC of the resampled data.
3. Repeat steps 1 and 2 many times until the distribution of AUC values converges.
4. From this bootstrap distribution, calculate the mean AUC and the desired percentile confidence interval.

This procedure is based on the pooled data bootstrapping approach, which has been shown to accurately estimate pharmacokinetic parameters (Mager and Goller 1998). As the experimental data used in this study comes from destructive sampling, dependence within a replicate across time is not considered; for experiments in which an individual is repeatedly sampled over time, the resamplings should take this dependence into account.

In step 2, the AUC is estimated by using the trapezoidal rule on the means of the replicates at each time point t_i where i=1,…,m and C_i is the average expression value for each time point after the resampling in step 1. Then, the AUC for this resampled data is given by

$$w_i = \begin{cases} \frac{1}{2}(t_2 - t_1), & i = 1 \\ \frac{1}{2}(t_{i+1} - t_{i-1}), & i = 2, \ldots, m-1 \\ \frac{1}{2}(t_m - t_{m-1}), & i = m \end{cases} \tag{1}$$

$$AUC = \sum_{i=1}^{m} w_i C_i . \tag{2}$$

In all of the results shown here (on both synthetic and real data), the number of resamplings performed (step 3) is 10000. This is well above the point by which the bootstrap distribution has converged, ensuring that the output generated in step 4 is stable.

To calculate the baseline, studies using baseline methods 1 and 2 (considering only the first time point or the first and final time points) likely will not have enough points to allow for a relatively smooth bootstrap distribution. As fewer data points are available, discreteness in the bootstrap distribution leads to biased confidence intervals (Bonate 1998). Thus, in these cases, the confidence interval is determined by assuming a normal distribution for the areas and using Bailer's method to calculate the variance. Bailer's method is a direct calculation of the variance of the AUC, where σ_i is the standard deviation of the expression values at time *i* for a given measure and r_i is the number of replicates (Bailer 1988).

$$\sigma^2_{AUC} = \sum_{i=1}^{m} w_i^2 \left[\frac{\sigma_i^2}{r_i} \right] \tag{3}$$

The above formulas for calculating AUC and its variance are taken over the interval [0, t_m]. However, if a response is known not to be resolved by t_m, these calculations can be further extrapolated. For instance, in a single bolus dose experiment, drug concentration can often be assumed to exponentially decay after t_m; applying this knowledge allows for the calculation of AUC over the interval [0, ∞], which may be a more appropriate metric in this particular case (Yuan 1993). In general, for gene expression and other types of data, such a relationship cannot be assumed, so the more conservative estimate over [0, t_m] should be used.

Calculating the AUC for biphasic responses

Gene expression data often contains biphasic or multiphasic responses. For instance, consider the data shown in Fig. 1. In the first gene, there is early down-regulation followed by late up-regulation before a return to baseline at the final time point. If the AUC and baseline estimates for these data were calculated as described above, they would be approximately the same because the deviations in gene expression caused by the up-regulation and down-regulation are of similar magnitudes. This would result in this gene not being selected as having a significant change from baseline.

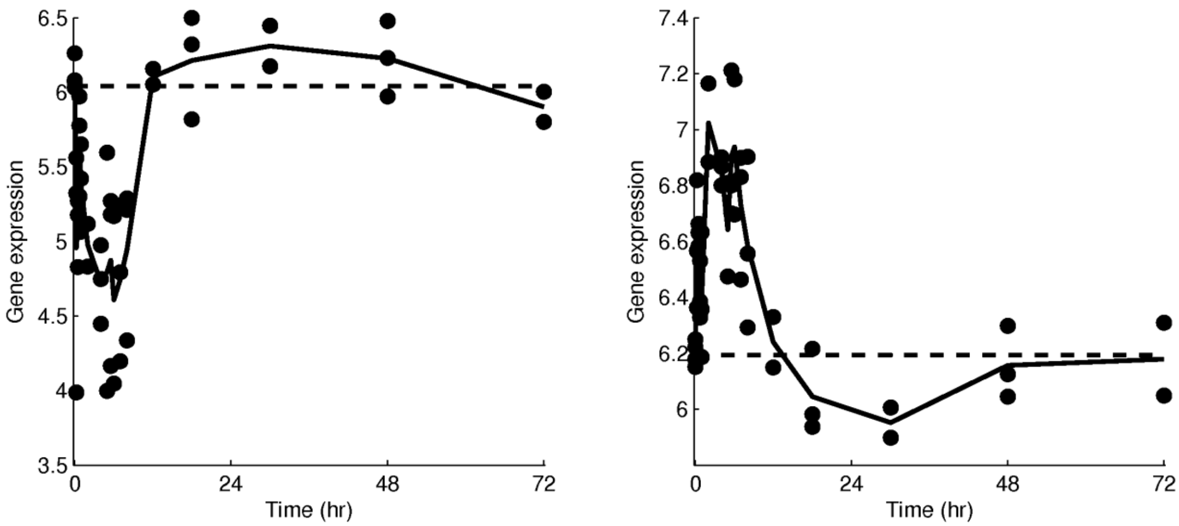

Fig. 1: Gene expression values for two genes are shown along with the average value at each time point (solid lines) and their baseline values (dashed lines). Left: The Hal gene shows acute downregulation followed by a rebound above baseline. Right: The Nudt4 gene shows a similar but inverted biphasic expression profile. By calculating the up-area and down-area separately, and then adding their absolute values, a more complete measure of the deviation from baseline is obtained.

This is seldom a problem in traditional applications of AUC in pharmacology because typically drug responses only go either up or down and then return to baseline. Occasionally, the occurrence of rebound is a complication in pharmacodynamics (Sharma, Ebling et al. 1998). But for gene expression data, up-regulation and down-regulation often occur sequentially, producing biphasic profiles. Thus, the trapezoidal rule is modified such that areas cannot be negative. In other words, the area above baseline is added to the area below baseline, and both of those areas are positive numbers. Algorithmically, this is accomplished by finding the point at which an interpolating line between two time points crosses the baseline value, and then calculating the area separately for the parts on the left and right of that point. This results in a large overall AUC for the genes in Fig. 1.

Selection of large AUCs

After calculating the baseline, the AUC, and their confidence intervals, these values are compared to determine if there is a significant difference between the baseline and the AUC. This is done by testing if the AUC confidence interval overlaps the baseline confidence interval. For all of the results shown in this paper, confidence intervals of 80% were used.

Detection of up/down/biphasic AUCs

Because the proposed algorithm for calculating AUC accounts for both positive and negative deviations from baseline, it is possible to distinguish between up-regulation, down-regulation, and biphasic responses. In the calculation of AUC, positive and negative contributions to AUC are tabulated separately. Then, the ratio between these AUC components and baseline variability is calculated, in keeping with the overarching idea that the magnitude of change relative to baseline is of interest. In the case when there is a response only in one direction followed by a return to baseline, then only one of the positive and negative AUCs will be large. If there are both large positive and large negative AUCs, this indicates a biphasic response.

After the significantly large AUCs are found as described above, biphasic responses are identified as those where the positive area is within 50% of the negative area and vice versa. Then, the remaining responses are classified as either up-regulation or down-regulation depending on whether the positive AUC or negative AUC is greater.

Synthetic data

Indirect response (IDR) models are widely used to represent inhibitory and stimulatory effects where the direct action of drugs is on a production or loss process (Krzyzanski and Jusko 1998; Woo, Pawaskar et al. 2009). Synthetic data were generated based on IDR model III, which models the indirect stimulatory effect *R* of a drug on k_{in} with the drug having an exponentially decaying concentration *C(t)*.

$$\frac{dR}{dt} = k_{in} \cdot \left(1 + \frac{S_{max} \cdot C(t)}{SC_{50} + C(t)}\right) - k_{out} \cdot R \tag{5}$$

$$C(t) = D / V \cdot e^{-(CL/V) \cdot t} \tag{6}$$

In these formulas, $R_0 = k_{in}/k_{out}$ is the initial response, S_{max} is the maximum effect, SC_{50} is the drug concentration producing 50% of the maximum effect, *D* is the intravenous dose, *V* is the volume, and *CL* is the clearance. For the synthetic data used in this study, parameter values of $R_0 = 50$, $CL = 2.5$, $D =$ 10,000, $V = 4$, $S_{max} = 5$, $SC_{50} = 4$, $k_{out} = 0.4$, and $k_{in} = 20$ were used, as in (Woo, Pawaskar et al. 2009). *R(t)* is sampled at the same time points as in the experimental data described below: 0, 0.25, 0.5, 0.75, 1, 2, 4, 5, 5.5, 6, 7, 8, 12, 18, 30, 48, and 72 hours. Qualitatively, this model shows an early up-regulation followed by a return to baseline. Therefore, the baseline value can be estimated by the measurements at the first and last time points. After the profile was generated for a given set of parameters, Gaussian noise with with a constant relative standard deviation (RSD) was added to generate an ensemble of noisy profiles. Noise is added based on the RSD, rather than using noise with the same standard deviation despite the current value of the signal, to simulate biological data in which there is a relationship between mean value and noise.

Acute corticosteroid gene responses

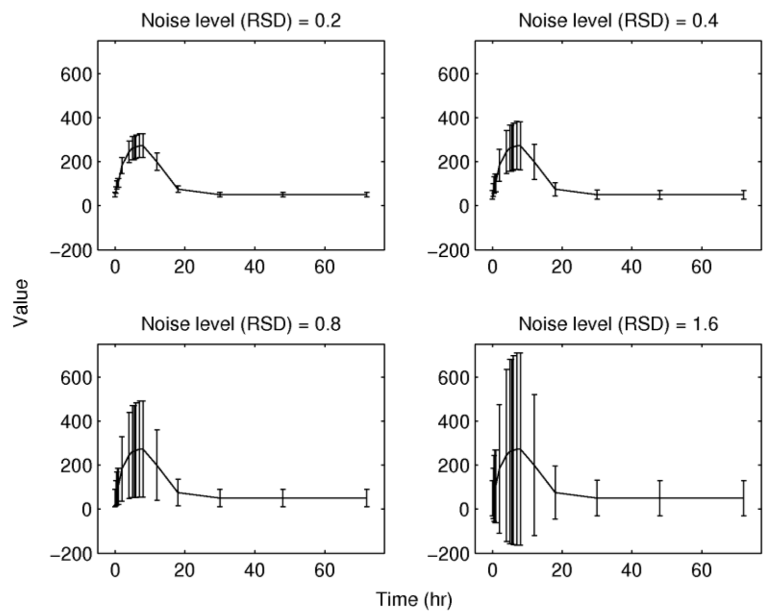

Fig. 2: Synthetic data generated for indirect response versus time profiles with four different noise levels at a single drug dose, D = 10,000.

In addition to the results obtained by using synthetic data in various conditions, the proposed algorithm was also run on real gene expression data to assess its ability to select biologically relevant genes. The data contains the transcriptional response of rat liver to a bolus dose of 50 mg/kg methylprednisolone (MPL) (Jin, Almon et al. 2003). It is available in the Gene Expression Omnibus (GEO) database (Barrett, Troup et al. 2009) with accession number GDS253. Forty-three male adrenalectomized Wistar rats were sacrificed at 16 time points: 0.25, 0.5, 0.75, 1, 2, 4, 5, 5.5, 6, 7, 8, 12, 18, 30, 48, and 72 hours after dosing; in addition, 4 more rats were sacrificed at 0 hours as a control group. Isolated RNA from each rat liver was hybridized to Affymetrix Rat Genome U34A microarrays to measure the expression values of 8,799 probe sets. This dataset generally consists of genes that have acute responses to the drug and eventually return to their original expression values. So, like for the synthetic data, the baseline value is estimated by taking the values of the replicates at the first and last time points. A confidence interval of 80% is used as the cutoff for determining significant AUC values.

Results

Synthetic data

Here 5000 synthetic profiles were generated based on the IDR model and 5000 flat profiles were created to simulate data with no time dependence. Gaussian noise, with RSD equal to 0.2, 0.4, 0.8, and 1.6, was added to the profiles to create four synthetic datasets with varying levels of noise, Fig. 2. To test the ability of the proposed algorithm to distinguish between these two classes, a receiver operating characteristic (ROC) curve was generated by varying the confidence interval cutoff from 0 to 1 and calculating the number of false positives and true positives in each scenario, as shown in Fig. 3. The dotted diagonal line represents the curve that would appear in the case of random guessing. This procedure was repeated for RSDs of 0.2, 0.4, 0.8, and 1.6. As the noise increases, the ROC curves move closer to the dashed line, but even at the highest noise level, the proposed algorithm still can somewhat distinguish between the two classes.

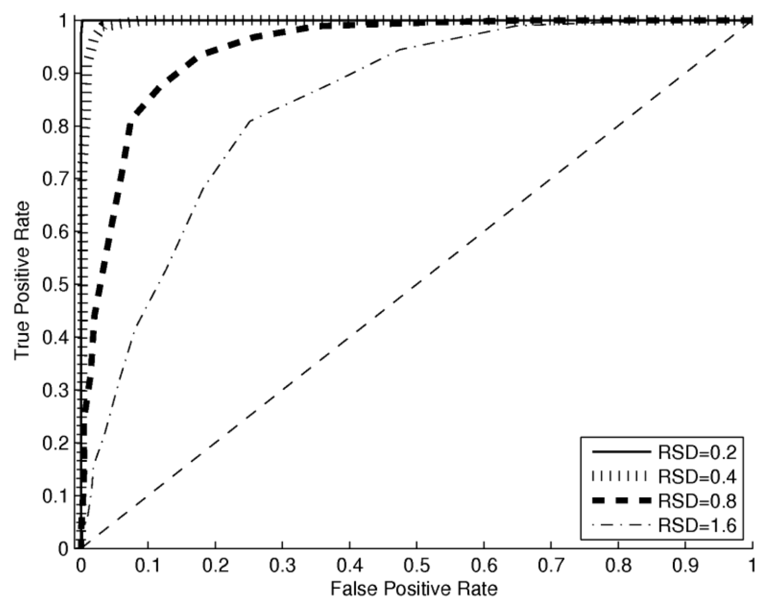

Fig. 3: ROC curve showing the ability of the proposed algorithm to accurately discover genes with an underlying time-dependent pattern with random noise at several relative standard deviations (RSD). RSD=0.2, AUROC=1; RSD=0.4, AUROC=0.995; RSD=0.8, AUROC=0.944; RSD=1.6, AUROC=0.839.

In these simulations, the drug dose was kept constant at D = 10,000. To investigate what happens as the drug concentration is changed, synthetic data were generated for values of D ranging from 10 to 1,000,000 as shown in Fig. 4. These profiles were run through the proposed selection algorithm to see what value of D is necessary to create an AUC that is significantly larger than baseline. Fig. 5 shows how the response of the proposed algorithm changes at these different initial drug concentrations. The black histogram represents all genes and the gray histogram represents selected genes. If the algorithm was perfectly able to determine that these genes all deviate from their baseline values, the two histograms would completely overlap.

Acute corticosteroid responses

When the algorithm is run on the acute corticosteroid dataset with a confidence interval cutoff of 80% and using the first and last time points to estimate the baseline, 345 genes have significantly large AUCs when compared with their baseline values. Biphasic responses are detected in many of these significant genes. The 33 genes that have both large positive and negative AUCs are shown in Table I along with the positive and negative contributions to AUC for each gene. Of the non-biphasic genes with significantly large AUCs, 139 are up-regulated and 173 are down-regulated based on the magnitude of their positive and negative AUCs. A full list of these genes, along with their positive and negative AUCs, is available as supplementary material.

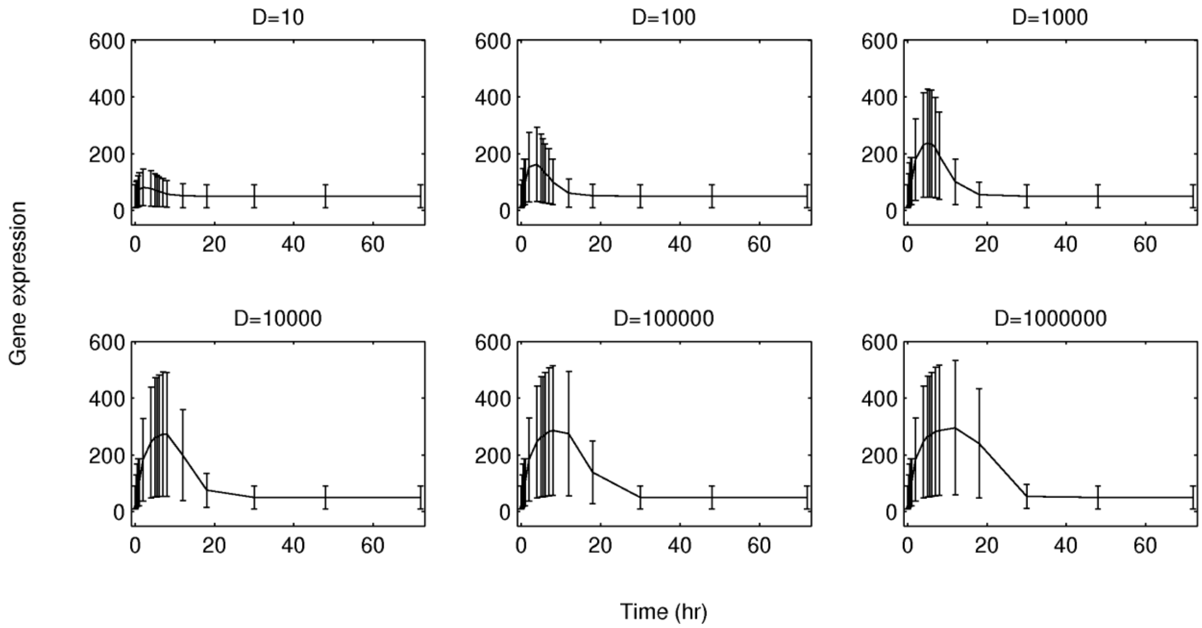

Fig. 4: Synthetic data generated for indirect response versus time profiles with 6 different dose (D) levels, all with the same amount of noise, RSD=0.8.

Discussion

Fig. 5 shows how the proposed method is able to discover significant AUCs over baseline values for a range of drug doses. At very low drug doses, it is impossible to distinguish between the two classes. As *D* increases, the algorithm performs better, but is still misses many genes. Interestingly, the significant genes are only slightly biased towards having higher than normal AUCs. This is because of the influence of the baseline value on the result of the significance test. If the measurements at the first and last time point are very close together, then the baseline estimate will have a very tight confidence interval. In these cases, the AUC estimate can be lower and still be called significant. But when there is a substantial uncertainty in the baseline value, only a very high AUC will be sufficiently greater than the baseline confidence interval. The critical importance of the baseline estimate on the output of this algorithm is a key difference between this AUC method and traditional gene expression significance tests. When considering a significance test designed to select patterns from high-dimensional data, it is important to consider the prevalence of false positives in the output (Storey and Tibshirani 2003). As a trivial example, running 1000 hypothesis tests at a p-value cutoff of 0.05 should lead to the rejection of the null hypothesis 50 times, given random data. The proposed algorithm does select a small number of genes from high-dimensional microarray data; however, when applying the algorithm to random data of the same size as the real experimental data, only ~10% as many genes are selected, suggesting that the vast majority of identified patterns do not arise by chance. Furthermore, biphasic responses are even less likely to arise by chance, with only ~5% as many biphasic genes detected in random data as in real

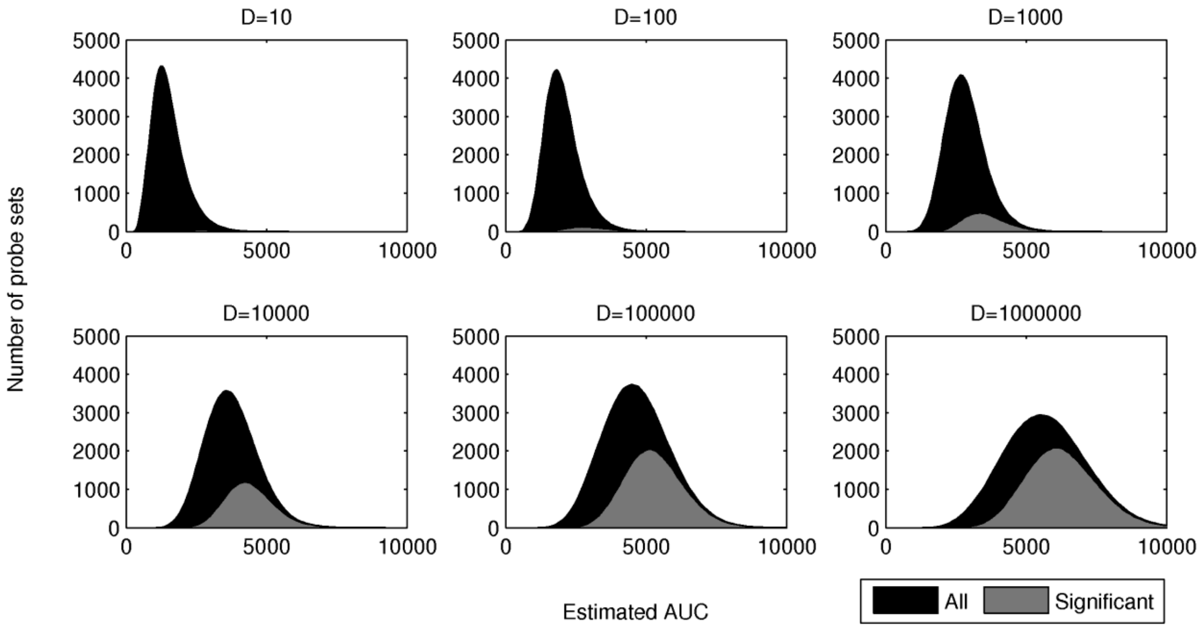

Fig. 5: Histograms of the AUCs of the synthetic data for various values of D given the same noise level RSD=0.8. The black histogram shows all of the synthetic genes, while the gray histogram shows only the significant synthetic genes. When the gray histogram is small, that means that few genes are selected by the algorithm. If all genes were selected, the black and gray histograms would overlap entirely.

data; this is likely because it is highly unlikely for equally large random perturbations to occur in both directions relative to baseline.

Bootstrapping is used to calculate the AUCs and their confidence intervals. While this is more computationally intensive and complex than a parametric calculation of the confidence interval, AUC values in general cannot always be accurately modeled by a normal distribution. To assuage its computational cost, the bootstrapping algorithm is parallelized (supplementary material, auc_bootstrap.m), which is important given that parallel processing capabilities are becoming uniquitous due to cheap multicore processors, even for researchers without access to dedicated computing facilities. Therefore, the increased complexity of nonparametric techniques such as bootstrapping is less of a concern than one might initially imagine.

While variability in AUC has previously been considered in the literature (Yuan 1993; Bonate 1998; Wolfsegger and Jaki 2005), this work considers this variability with respect to uncertainty in the baseline measurement which is of importance in any pharmacokinetic profile or pharmacodynamic response that does not have a baseline of exactly zero. The baseline has been incorporated into AUC calculation in the case when baseline measurements from a control group are available at all time points (Jaki, Wolfsegger et al. 2009), but the proposed method is flexible in that baseline is determined from any available experimental data even if it is just a single time point at *t*=0.

An advantage of the proposed algorithm is that it can automatically detect biphasic responses (such as those shown in Fig. 1) that have large positive and negative deviations from baseline. Table I lists the genes that were found to have large positive and negative deviations from baseline and Fig. 1 displays the time course plots of two of those selected genes. In the absence of the separate calculation of positive and negative AUCs, the AUC would be calculated as the difference between those two values. For the identified biphasic genes, these values are given in the last two columns of Table I. In each case, the positive and negative AUCs are sufficiently close that they would almost completely cancel out if they were not calculated separately. Because the genes in Table I were selected by calculating the AUCs relative to variable baselines, the positive and negative areas represent significant deviations from baseline. Without taking this variability into account, a signal that just oscillated above and below its baseline value due to random noise might be identified as having both positive and negative AUC values; but, since those oscillations would be small compared to the baseline variability, that signal would not be selected as having biphasic characteristics. Thus, the calculation of AUC compared to baseline can automatically detect biphasic responses. However, for general gene expression analysis, it does not perform as well as domain-specific algorithms such as EDGE (Leek, Monsen et al. 2006) for the task of identifying differentially expressed genes (data not shown). The proposed algorithm is unique in that is can search directly for biphasic responses, thus effectively performing some classification (up/down/biphasic) of the responses at the same time as testing for significant deviations in AUC relative to baseline.

Caution is needed when applying any AUC calculation method to incomplete pharmacodynamic data. Unlike pharmacokinetics which usually exhibits a monoexponential terminal phase allowing easy extrapolation to time infinity, the return phase for dynamic data to baseline is nearly, but not exactly linear (i.e. $\Delta R/\Delta t$ = constant) (Krzyzanski and Jusko 1998).

The proposed method of calculating the AUC relative to a variable baseline can be applied in discovering significant differentially expressed genes, as an alternative to existing methods, and also to determining specifically which genes are up-regulated, down-regulated, and biphasic. As with noncompartmental methods in pharmacokinetics, this type of analysis should have value in screening large data sets and as a preliminary step before application of specific models with use of more sophisticated regression software. Furthermore, its utility is not restricted to data from microarray experiments. The proposed method can characterize the AUC for any time series data with a non-constant baseline as is commonplace in pharmacodynamics. An implementation of this algorithm in MATLAB is available as supplementary material.

Acknowledgements

JDS and IPA acknowledge support from NIH Grant No. GM082974. RRA, DCD and WJJ acknowledge support from NIH Grants No. GM24211 and GM57980.

References

Bailer, A. J. (1988). "Testing for the equality of area under the curves when using destructive measurement techniques." J Pharmacokinet Biopharm **16**(3): 303-9.

Barrett, T., D. B. Troup, et al. (2009). "NCBI GEO: archive for high-throughput functional genomic data." Nucleic Acids Res **37**(Database issue): D885-90.

Bonate, P. L. (1998). "Coverage and precision of confidence intervals for area under the curve using parametric and non-parametric methods in a toxicokinetic experimental design." Pharm Res **15**(3): 405-10.

Ernst, J. and Z. Bar-Joseph (2006). "STEM: a tool for the analysis of short time series gene expression data." BMC Bioinformatics **7**: 191.

Jaki, T., M. J. Wolfsegger, et al. (2009). "Confidence intervals for ratios of AUCs in the case of serial sampling: a comparison of seven methods." Pharm Stat **8**(1): 12-24.

Jin, J. Y., R. R. Almon, et al. (2003). "Modeling of corticosteroid pharmacogenomics in rat liver using gene microarrays." J Pharmacol Exp Ther **307**(1): 93-109.

Krzyzanski, W. and W. J. Jusko (1998). "Characterization of pharmacodynamic recession slopes for direct and indirect response models." J Pharmacokinet Biopharm **26**(4): 409-36.

Krzyzanski, W. and W. J. Jusko (1998). "Integrated functions for four basic models of indirect pharmacodynamic response." J Pharm Sci **87**(1): 67-72.

Kwon, Y. (2001). Handbook of essential pharmacokinetics, pharmacodynamics, and drug metabolism for industrial scientists. New York, Kluwer Academic/Plemun Publishers.

Leek, J. T., E. Monsen, et al. (2006). "EDGE: extraction and analysis of differential gene expression." Bioinformatics **22**(4): 507-8.

Mager, H. and G. Goller (1998). "Resampling methods in sparse sampling situations in preclinical pharmacokinetic studies." J Pharm Sci **87**(3): 372-8.

Sharma, A., W. F. Ebling, et al. (1998). "Precursor-dependent indirect pharmacodynamic response model for tolerance and rebound phenomena." J Pharm Sci **87**(12): 1577-84.

Storey, J. D. and R. Tibshirani (2003). "Statistical methods for identifying differentially expressed genes in DNA microarrays." Methods Mol Biol **224**: 149-57.

Tusher, V. G., R. Tibshirani, et al. (2001). "Significance analysis of microarrays applied to the ionizing radiation response." Proc Natl Acad Sci U S A **98**(9): 5116-21.

Wolfsegger, M. J. (2007). "Establishing bioequivalence in serial sacrifice designs." J Pharmacokinet Pharmacodyn **34**(1): 103-13.

Wolfsegger, M. J. and T. Jaki (2005). "Estimation of AUC from 0 to infinity in serial sacrifice designs." J Pharmacokinet Pharmacodyn **32**(5-6): 757-66.

Woo, S., D. Pawaskar, et al. (2009). "Methods of utilizing baseline values for indirect response models." J Pharmacokinet Pharmacodyn **36**(5): 381-405.

Yuan, J. (1993). "Estimation of variance for AUC in animal studies." J Pharm Sci **82**(7): 761-3.

Tables

Table I: Positive and negative contributions to the AUC for genes with biphasic responses

Gene name	Accession no.	Positive AUC	Negative AUC
Hal	AB002393_g_at	8.72	11.81
Inhbe /// LOC497821	AF089825_at	5.26	7.77
Mosc2	AF095741_at	9.77	8.60
Mosc2	AF095741_g_at	5.85	6.82
Cyb5r3	D00636Poly_A_Site#1_s_at	5.64	10.73
Csda	D28557_s_at	8.41	6.05
Psmc4	D50695_at	5.32	9.81
Abcc2	D86086_s_at	7.95	7.29
Psma2	E03358cds_g_at	5.19	3.63
Eif2s1	J02646_at	3.44	6.58
Tpm4	J02780_at	5.66	4.45
Nr1d1	M25804_g_at	5.50	10.70
Agtr1a	M86912exon_g_at	3.57	6.53
Vhl	rc_AA799545_at	6.99	11.60
Rara	rc_AA799779_g_at	7.50	7.90
Serbp1	rc_AA800678_g_at	14.95	9.01
Akap12	rc_AA859966_i_at	12.22	8.77
Phb2	rc_AA875054_at	3.74	2.32
Pxmp2	rc_AA875639_at	8.35	6.35
Cct3	rc_AA891107_at	8.98	8.11
Gnpat	rc_AA894174_g_at	3.14	2.47
LOC363328	rc_AA894305_at	8.90	5.34
	rc_AA943892_at	3.69	4.55
Tcp1	rc_AI104500_at	7.60	3.96
Zadh2_predicted	rc_AI229637_at	5.52	5.78
Nudt4	rc_AI237535_s_at	4.54	7.33
Etfa	rc_AI639026_at	3.97	2.35
RGD1308373	U14746_at	5.47	10.72
Agt	U15211_g_at	5.97	6.86
Entpd1	U21718mRNA_at	4.74	8.07
Mybbp1a	U23146cds_s_at	8.32	9.81
Litaf	U75392_s_at	3.66	2.14
	X70223_at	5.68	8.20

Positive and negative AUC contributions for the 33 genes identified as having biphasic responses are listed. AUC values have units of log(intensity)*hr, where the intensity value is calculated from the microarray experiments.