

## A Bayesian approach to analyzing phenotype microarray data enables estimation of microbial growth parameters

Matthias Gerstgrasser<sup>\*¶</sup>, Sarah Nicholls<sup>†||</sup>, Michael Stout<sup>†\*\*</sup>,  
Katherine Smart<sup>‡††</sup>, Chris Powell<sup>‡‡‡</sup>, Theodore Kypraios<sup>§§</sup>  
and Dov Stekel<sup>¶¶</sup>

*\*Department of Computer Science, University of Oxford  
Parks Road, Oxford, OX1 3QD, UK*

*†School of Biosciences  
University of Nottingham  
Sutton Bonington, Leicestershire, LE12 5RD, UK*

*‡SAB Miller plc, Church Street West  
Woking, GU21 6HS, UK*

*§School of Mathematics, University of Nottingham  
Nottingham, NG7 2RD, UK*

*¶matthias.gerstgrasser.net  
||sarah.nicholls@nottingham.ac.uk  
\*\*michael.stout@nottingham.ac.uk*

*††katherine.smart@sabmiller.com  
‡‡chris.powell@nottingham.ac.uk*

*§§theodore.kypraios@nottingham.ac.uk  
¶¶dov.stekel@nottingham.ac.uk*

Received 1 May 2015

Revised 3 September 2015

Accepted 18 October 2015

Published 13 January 2016

Biolog phenotype microarrays (PMs) enable simultaneous, high throughput analysis of cell cultures in different environments. The output is high-density time-course data showing redox curves (approximating growth) for each experimental condition. The software provided with the Omnilog incubator/reader summarizes each time-course as a single datum, so most of the information is not used. However, the time courses can be extremely varied and often contain detailed qualitative (shape of curve) and quantitative (values of parameters) information. We present a novel, Bayesian approach to estimating parameters from Phenotype Microarray data, fitting growth models using Markov Chain Monte Carlo (MCMC) methods to enable high throughput estimation of important information, including length of lag phase, maximal “growth” rate and maximum output. We find that the Baranyi model for microbial growth is useful for fitting Biolog data. Moreover, we introduce a new growth model that allows for diauxic growth with a lag phase, which is particularly useful where Phenotype Microarrays have been applied to cells grown in complex mixtures of substrates, for example in industrial or

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC-BY) License. Further distribution of this work is permitted, provided the original work is properly cited.

biotechnological applications, such as worts in brewing. Our approach provides more useful information from Biolog data than existing, competing methods, and allows for valuable comparisons between data series and across different models.

*Keywords:* Biolog; growth model; diauxic; lag phase; Bayesian statistics; phenotype microarrays.

## 1. Background

Biolog Phenotype Microarrays (PMs) are unique patented commercial products for assessment of cellular respiration of prokaryotic and eukaryotic cells in a wide range of conditions, including metabolism using different carbon, nitrogen, phosphorous and sulphur sources, as well as osmotic, pH, antimicrobial and metal ion stresses. The PMs work by the reduction of a colorless tetrazolium dye in the growth media to a purple formazan by electrons from the NADH produced during cellular respiration.<sup>1</sup> For microbial PMs, 1920 different phenotypes per organism (including controls) can be assessed simultaneously by using the full set of 20 different 96 well plates. As the assays are performed for 24 h or longer, the output is high density time-course data for each well (growth condition), showing a measurement of the quantity of dye reduced. A typical experiment can contain as many as 450,000 data points, making the output especially suitable for mathematical and statistical modelling. The PM platform is flexible, allowing users to construct their own assays using plates, growth media and tetrazolium dyes. Thus, they have proven extremely flexible in terms of the experiments that can be carried out with them.<sup>2</sup> Among other uses, Biolog Omnilog PM technology can be used in research aimed at understanding and controlling the performance of biotechnological processes, through analysis of microbial metabolism in conditions relevant to industrial fermentations.<sup>3,4</sup>

However, while PM technology can generate vast amounts of information about microbial growth in the form of time-series data, its usefulness is limited by a lack of robust, easy to use and flexible data analysis tools for such data. Often, the data (which in its raw form comprises up to several hundred data points) is being reduced to either a binary “growth/no-growth” distinction or, at best, a single datum, such as maximum signal reading, average signal height reading or area under curve (AUC).<sup>5,6</sup> Clearly, a lot of information is lost this way.

Biolog’s own software allows rudimentary data analysis mainly focused at directly comparing two different strains grown on PM plate types. As detailed in Bochner *et al.*<sup>5</sup> — cf. Box 1 therein — this consists of plotting curves from two strains (generally a mutant and a control strain) against each other and highlighting differences above a certain threshold. This results in effectively a ternary distinction between either of the strains showing higher respiration or there being no significant difference between them. Alternatively the software can give a numerical output of the difference in respiration rates between the two strains compared, as used for instance in Ref. 7.

Whether average readout, AUC or endpoint is used, the use of a single value to represent a time-series comprising up to several hundred data points entails losing valuable information about the shape of the underlying curve. For instance, the

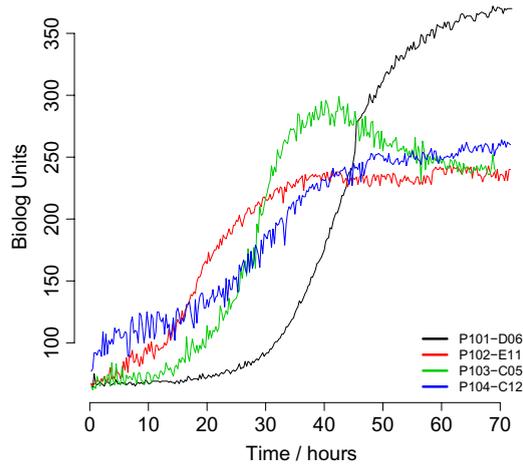


Fig. 1. Respiration data taken from four different phenotypes in the datasets. Plot shows coloration (Biolog units) versus time. All four curves have very similar AUC, yet it is clear that they are qualitatively different, with different lengths of lag phase and maximal growth rates.

curves in Fig. 1 all have very similar average readout and AUC. Yet clearly they are qualitatively rather different, with different lengths of lag phases, maximal growth rates and carrying capacities. This is the problem at the heart of Biolog data analysis that we seek to address.

Some previous research has been carried out in this area, but has been limited in its scope. For instance, Ref. 8 focuses on visualizing the raw Biolog data without any parameter estimation. A more recent approach published in Ref. 30 provides both visualization as well as parameter estimation using the **grofit** R package. This package uses nonlinear least squares regression to fit Gompertz and Richards models to growth curves, but also provides a model-free spline fit.<sup>10</sup> The logistic growth model has also been found to be effective in fitting Biolog data and has been used to facilitate normalization and data comparison.<sup>11</sup>

In this paper, we describe a method to extract further information from these curves. We employ a Bayesian approach using Markov Chain Monte Carlo (MCMC) techniques to sample from the posterior distributions of the parameters of several different growth models fitted to given Biolog data. Such an approach not only offers robust estimates of both best-fit model parameters but also model-independent characteristics such as lag time, maximal growth rate, and maximal carrying capacity. We have tested these methods using data from customized PM plates, developed to assess the potential of PM technology in the differentiation of 100 proprietary brewing yeast strains in different worts. The focus of the paper is on the modeling and methodology and these data are brought as relevant examples.

We anticipate that these ideas have the potential to transform the capacity of research groups to obtain useful and meaningful information from Biolog Phenotype Array data.

## 2. Methods

### 2.1. Modeling aspects

A number of different growth models have been proposed in the literature.<sup>12</sup> These have different statistical properties<sup>13</sup> and it can be argued that there are limitations of the interpretations of these models.<sup>14</sup> These models are generally derived in the form of an ordinary differential equation or as a system of ODEs, though many of them also have a closed-form expression. We will focus on two models in particular, described in the sections below. As the Biolog Omnilog PM machines measure respiration rather than growth, it is not *a priori* clear that any of the traditional growth models will be able to accurately fit the data produced in PM experiments. However, empirically we found that these models provide useful results.

### 2.2. The Baranyi model

We chose to focus mainly on the model developed by Baranyi and Roberts. This was first introduced in Ref. 15, and discussed in more detail for instance in Refs. 16 and 17. The model is based on the Richards model,<sup>18</sup> but introduces another inhibition term to model the lag phase. The inclusion of a lag phase is important, as growth and metabolism of microorganisms in fresh media typically results in an initial period of delayed activity. Consequently, much of the data we analyzed includes a distinct lag phase; this can be seen in the example data shown in Fig. 1. Thus, models that do not include a lag phase (e.g. logistic growth), or models where there is insufficient flexibility over the shape and time of the lag phase (e.g. Gompertz models<sup>12</sup>) do not perform well. The form of the model we used is as in Ref. 16:

$$\frac{d}{dt}y = r \cdot y \cdot u(y) \cdot \alpha(t), \quad (1)$$

where  $\alpha(t)$  accounts for the inhibition at the beginning of growth. For an isothermal batch culture environment (e.g. as provided by Omnilog), the authors suggest setting

$$\alpha(t) = \frac{q_0}{q_0 + \exp(-\nu \cdot t)}. \quad (2)$$

With  $u(y)$  as in the Richard's model this gives a closed-form expression as follows:

$$\log(y(t)) = \log(y_0) + r \cdot A(t) - \frac{1}{m} \log \left( 1 + \frac{e^{mrA(t)} - 1}{e^{m(\log(y_{\max}) - \log(y_0))}} \right), \quad (3)$$

where

$$A(t) = \int_0^t \alpha(s) ds = t + \frac{1}{\nu} \log \left( \frac{e^{-\nu t} + q_0}{1 + q_0} \right). \quad (4)$$

In addition to the four parameters ( $y_0, y_{\max}, r, m$ ) of the Richards model, this introduces another two parameters,  $\nu$  and  $q_0$ , that control the length and shape of the lag phase. This term is motivated biologically:  $q_0$  is to be taken as the initial

physiological state of the cells, while  $\nu$  gives the rate at which they adapt to their new environment. This gives the full Baranyi model a high degree of freedom in accommodating a wide range of growth curve data; as others have noted<sup>13</sup> the Baranyi model performs very well in fitting empirical growth curve data.

There are two things we would like to point out. First, the form of  $\alpha(t)$  makes it entirely independent of the rest of the Baranyi model. That is, it is straightforward to incorporate this lag term into other models, and we have done so for instance with a simple diauxic model which we discuss below. Second, for any given lag phase length  $\lambda$ , there is an infinite number of combinations of  $q_0$  and  $\nu$  that satisfy the formula for  $\alpha$  given above. We have found it more convenient to parameterize the Baranyi model (and other models using  $\alpha(t)$ ) with  $\lambda$  and  $\nu$  rather than  $q_0$  and  $\nu$ . We then derive  $q_0$  by

$$q_0 = \frac{1}{e^{\nu\lambda} - 1}, \quad (5)$$

for use in the closed-form expression of the Baranyi model respectively  $\alpha(t)$ . In this form, one parameter controls the duration of the lag phase itself, whereas  $\nu$  controls its shape. Figure 2(b) shows  $A(t)$  for a fixed  $\lambda$  and varying values of  $\nu$ . Higher values of  $\nu$  make the lag phase more pronounced, whereas lower values make the effect more subtle (in the figure, the topmost curve is the one with the lowest value for  $\nu$ ).

Baranyi suggested a slight simplification of the original model, setting  $m = 1$  and  $\nu = r$ <sup>19</sup>. The resulting model is simpler than the original one and still fits most of the data we have seen; On the other hand, the original model is more flexible in fitting less typical growth curves. We thus adopted a two-fold strategy and used both the original and the simplified versions of the model, as well as two versions that incorporate either one of the proposed modifications.

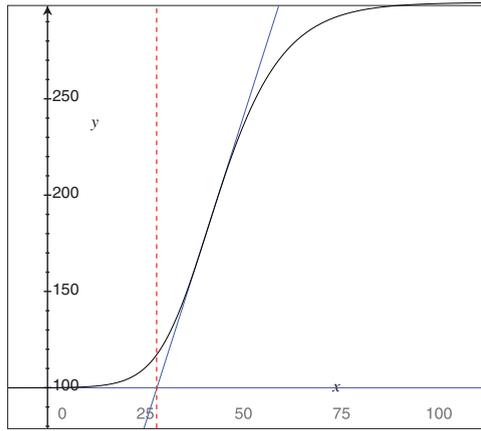
### 2.3. Diauxic growth

As will be seen below, the example data we analyze (from worts), typical of growth on complex mixtures of substrates, includes curves that show a clear diauxic effect. We modeled this using a simple diauxic growth model based on Monod-type substrate inhibition terms. As a starting point, we used the following model, in ODE form:

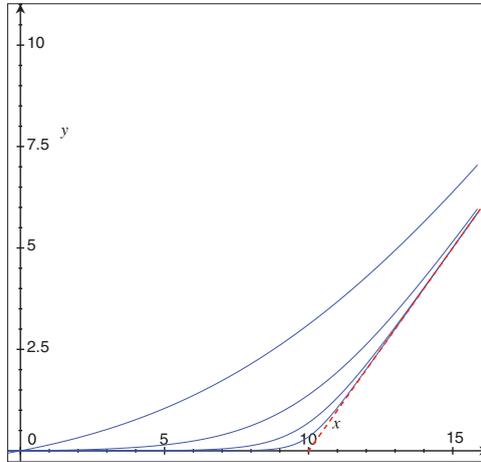
$$\frac{dy}{dt} = r_1 s_1 y + \frac{k}{k + s_1} r_2 s_2 y, \quad (6)$$

$$\frac{ds_1}{dt} = -r_1 s_1 y, \quad \frac{ds_2}{dt} = -\frac{k}{k + s_1} r_2 s_2 y. \quad (7)$$

While  $s_1$  is large, this essentially gives logistic growth on substrate 1. As substrate 1 is being used up, the inhibition term  $\frac{k}{k+s_1}$  increases and logistic growth on substrate 2 starts. Smaller values of the inhibition constant  $k$  give a more pronounced diauxic



(a)



(b)

Fig. 2. Numerical derivation of a lag phase length from a Baranyi curve. (a) The calculated lag is 27.87 whereas the model lag parameter is 30. (Other model parameters:  $y_0 = 100, y_{\max} = 300, r = 0.1, \nu = 0.2, m = 1$ ) and (b)  $A(t)$  for  $\lambda = 10$  and  $\nu = 0.2, 0.5, 1, 2$ . Dotted red line shows  $y = t - \lambda$ .

effect due to the stronger influence of  $s_1$  on the inhibition term.  $\frac{k}{k+s_1}$  is a hyperbolic inhibition term that is supposed to model the inhibitory effects happening within the cells as  $s_1$  is being exhausted.

We amended this model to include a Baranyi-like lag phase term. In its non-integral form  $\alpha(t)$  this transfers to our ODE model in a straightforward fashion, giving the following for our diauxic growth model:

$$\frac{dy}{dt} = \alpha(t) \left( r_1 s_1 y + \frac{k}{k + s_1} r_2 s_2 y \right), \tag{8}$$

$$\frac{ds_1}{dt} = -\alpha(t)(r_1s_1y), \quad \frac{ds_2}{dt} = -\alpha(t)\left(\frac{k}{k+s_1}r_2s_2y\right). \quad (9)$$

Arguably, this model is a rather simplified view of the biological processes happening but has empirically proven to be useful. A number of other models for diauxic growth have been proposed.<sup>20,21</sup> Our model is simpler (for example Kompala *et al.*'s 1986 model has five parameters per substrate, ours has eight in total), but still sufficiently flexible to be able to describe virtually all of the diauxic growth curves we have seen.

#### 2.4. Gompertz model

For comparison purposes, we also fitted a Gompertz model<sup>12</sup> to the data. The Gompertz model does not include a lag phase, so it is useful to determine when the Baranyi model is helpful in identifying lag phase lengths.

#### 2.5. Parameter estimation

We adopted a Bayesian approach<sup>22</sup> to infer model parameters from the time-series data recorded by Biolog machines. In particular we used a variant of the Metropolis-Hastings algorithm<sup>23</sup> to sample from the posterior distribution of these parameters. This algorithm starts with an initial state vector  $\theta^{(0)}$ . At each iteration, a candidate state vector  $\theta'$  is generated by drawing from a proposal distribution  $q(\cdot|\theta^i)$ . With probability  $\alpha(\theta_i, \theta')$  that move is accepted, where

$$\alpha(\theta_i, \theta') = \min\left(1, \frac{p(\theta')q(\theta^i|\theta')}{p(\theta^i)q(\theta'|\theta^i)}\right). \quad (10)$$

If accepted, we set  $\theta^{i+1} = \theta'$ , otherwise set  $\theta^{i+1} = \theta^i$ . The  $\theta^i$  form a Markov chain, and the stationary distribution of this chain is the desired posterior distribution  $p(\theta)$  irrespective of the proposal distribution  $q(\cdot|\cdot)$  in Refs. 24 or 25.

A number of modifications and amendments to the original Metropolis-Hastings algorithm have been proposed to better explore the target distribution or to improve the algorithm's rate of convergence. In particular, significant attention has been drawn to the construction of *adaptive algorithms*, i.e. algorithms which do not require the specification of the tuning variance of the proposal distribution  $q(\cdot|\cdot)$ . One of the most commonly used algorithms is the Adaptive Metropolis (AM) algorithm<sup>26</sup> that has been discussed in detail in the literature<sup>27-30</sup>.

The use of an adaptive algorithm is important for the analysis of high throughput PM data because of the large number of data sets being analyzed. Each PM plate contains 96 wells, and with an experiment of 100 plates this leads to 9600 separate curves. It is not possible to manually tune the parameters for each curve so an automated, adaptive approach is necessary. We implemented an AM algorithm

with global scaling as described in algorithm 4 in Andrieu and Thoms 2008 as follows<sup>2</sup>:

- (1) For an initial segment of the chain ( $i < i_0$ , for some sensible choice of  $i_0$ ) perform a “Random-Walk Metropolis with global scaling” with proposal covariance matrix  $\lambda_i \Sigma_0$ , with  $\Sigma_0$  an initial “best guess” of the true covariance matrix.
- (2) For the remainder of the chain, do as above but use  $\lambda_i \Sigma_i$  as the covariance matrix for the proposal distribution, where  $\Sigma_i$  is the sample covariance matrix of the previous history of the chain, and  $\lambda_i$  is a varying scaling factor. We update  $\Sigma_i$  iteratively.

We used  $i_0 = 2500$  (chosen empirically). For  $i < i_0$  we updated  $\Sigma_i$  and  $\lambda_i$  only when we accept a move, as suggested in Haario *et al.* 2001.<sup>26</sup> We reset  $\lambda_i$  to its original value of  $2.4^2 / \dim(\theta)$  in iteration  $i_0$ , where  $\theta$  is the parameter vector.<sup>2</sup> We also added  $\epsilon I$  to  $\Sigma_i$  in each step to keep it from becoming singular for some small  $\epsilon$ <sup>26</sup>. If ever it still became singular due to rounding errors, we added  $\epsilon I$  to it until it was not singular anymore. We did not use any form of thinning, always discarded the initial 50,000 iterations as burn-in, and in production runs used 500,000 iterations in total.

We used uninformative or uniform priors on suitable regions, and utilized simple heuristics to give rough estimates of the initial parameter vector for the AM algorithm (see Appendix for details). These heuristics proved to be effective in that all Markov chains converged to a stationary distribution (see Sec. 2.6 on Performance below) after appropriate burn-in. Alternative approaches could be to use an iterative least-squares approach to obtain initial parameter values near to the stationary distribution.<sup>11</sup> To calculate the likelihood function, we assumed normally i.i.d. measurement errors, which we heuristically estimated from the raw data as detailed in the appendix.

## 2.6. Model choice

In addition to inferring parameters from them, we used the Deviance Information Criterion.<sup>31,32</sup> This is defined as follows:

$$DIC = \bar{D} + p_V, \tag{11}$$

where  $D$  is the deviance defined by

$$D(\theta) = -2 \log(p(x|\theta)) \tag{12}$$

and  $\bar{D}$  is the mean of this deviance. The model complexity  $p_V$  is given as  $p_V = \text{var}(D)/2$ . An alternate definition of DIC uses  $p_D = \bar{D} - D(\bar{\theta})$  to account for model complexity. We have found however that this is highly dependent on the quality of the estimation of the posterior. In some cases where our posterior sample was not a good estimate, we saw the  $p_D$  term dropping to artificially low (often negative) values leading to a bad, artificially low estimate of DIC. While it is possible to check that the posterior sample is reasonable before calculating DIC this way, we

expect that in any high-throughput environment there will be individual data that slip through such checks. Thus, we have found that using  $p_V$  was in practice more suitable for our purposes. For comparison purposes we have also calculated BIC for all models,<sup>33</sup> using the highest likelihood observed in the posterior sample for the BIC estimate.

We used the DIC to choose which model to use to derive estimated parameters for each well. Furthermore, we were looking to get an indication as to whether our models provide a reasonable fit at all. To this end we compared them to a simple “dummy” model defined by  $y(t) = c$  for a constant parameter  $c$ . This is non-informative, and similar to a classical  $H_0$  model. A comparison via DIC or BIC to the constant model is meant to show whether the model fits the data at all.

## 2.7. Parameter comparisons between models

One potential issue with fitting a number of different models to data is that parameters of one model do not necessarily relate easily to those of another. Therefore, in addition to the model-specific parameters we used model-independent measurements of three key growth characteristics, which we derived numerically from the fitted curves. This is to allow for simple quantitative comparisons between wells that show qualitatively different behavior.

In the following paragraphs, we will take  $y(t; \theta)$  to mean the signal level predicted by our model and the parameter vector  $\theta$  at time  $t$ . First, for  $A$ , the maximum coloration change achieved, we simply used  $y(t_{\text{last}}; \theta) - y(0; \theta)$ , that is, the absolute increase in coloration our model predicts over the period of time that was recorded in the experiment. It may be argued that for all the models we are using we could also compute a similar  $A$  from the  $y_0$  and  $y_{\text{max}}$  (respectively  $y_0$  and  $s_1, s_2$ ). However, we found that if the recorded respiration data stops before a maximum is attained, the maximum or substrate level parameters are effectively inestimable, and so would  $A$  be with such a definition. Note that if a maximum is effectively reached within the experiment record, these two definitions will for practical purposes be equivalent.

For  $\mu_{\text{max}}$  we took the steepest slope of  $y(t; \theta)$  (again within the experiment period), as derived numerically from the fitted curve. This is not a transformation of the model’s rate parameter(s) alone, but is model-independent and coincides with a more natural definition of the maximal respiration rate.

For the lag time  $L$  an essentially model-independent definition is slightly trickier, and a number of possible definitions could be used.<sup>14</sup> We define  $L$  to be the  $t$ -coordinate of the intersection of the tangent at the steepest point of  $y(t; \theta)$  with the line  $y \equiv y(0, \theta)$ .<sup>13</sup> For  $y_0 \sim y_{\text{max}}$  this will approximate the lag parameter in  $\alpha(t)$  almost exactly. Figure 2(a) illustrates our definition of the lag phase length. The tangent of the steepest point is shown in blue, and the derived lag length  $L$  in red. We note that for  $y_0 \ll y_{\text{max}}$  this estimate will differ substantially from the lag parameter  $\lambda$ . However, so long as the relative difference between initial and final cell concentration is constant across a dataset, our estimate of  $L$  will still be comparable between wells and, crucially, between different models.

## 2.8. Identifying presence and absence of growth

In order to identify whether a given well exhibits significant levels of growth at all, we compared the maximum coloration attained  $A$  to a 95% quantile of the same parameter for a control well present on each plate.

## 2.9. Data preprocessing

In a small number of cases we observed anomalous behavior where coloration actually decreases significantly after attaining a maximum. The causes of this behavior are as yet unknown (in principle the reduction of the tetrazolium dye should be a one-way process), but we still want to be able to extract meaningful information from the data. We used a simple heuristic to remove the aberrant parts of the data. More precisely, we looked at the maximum  $y'_{\max}$  (attained at  $t_{\max}$ ) of a smoothed curve as above (using a 9-point window). We then removed the tail of the data series if in any interval  $[t_{\max}, t]$ ,  $t > t_{\max}$  at least 90% of data points (of the unsmoothed curve) were at least 0.5 standard deviations (using the numerically estimated measurement error) below  $y_{\max}$ . We would then remove the tail of the data series after the least such  $t$ . However, we always left at least the initial 40 data points.

# 3. Results and Discussion

## 3.1. Our models fit the data

The model was applied to 40 Biolog arrays comprising a total of 3840 time courses. Of these, 2989 exhibited significant growth compared to a known control well. According to DIC 598 of these wells were fitted best by the Baranyi model, 2382 by the diauxic model, 9 by the Gompertz model and none by the constant dummy model. The remaining 851 wells exhibited no significant growth, and model fit between the Baranyi, diauxic and Gompertz models was almost arbitrary.

According to BIC the picture is similar, with a slight bias against the diauxic model (2160 diauxic cases with growth) and toward simpler models (793 Baranyi, 36 Gompertz). Again no wells were fitted by the constant dummy model.

Furthermore, in all cases in which the Gompertz model was preferred over Baranyi or diauxic models, DIC scores of these models were very close to each other (within 1% of absolute values). The converse was not the case. Best-fit models usually achieved a DIC score in the range of 200–400. For cases where the Gompertz model scored best, the mean difference between the Gompertz model and the second-best model was 1.7; Conversely, where the Gompertz model did not score best, the mean difference in BIC score between the best model and the Gompertz model was 658.

This demonstrates that the Baranyi model, and the diauxic model with Baranyi lag phase, are effective tools for the analysis and interpretation of Biolog PM data.

Example fits where the Baranyi model and the diauxic model with Baranyi lag fit best are shown in Figs. 3(a) and 3(b), respectively. Plotted also is a fitted Gompertz model, demonstrating the relevance of the models we have used. In these cases, the

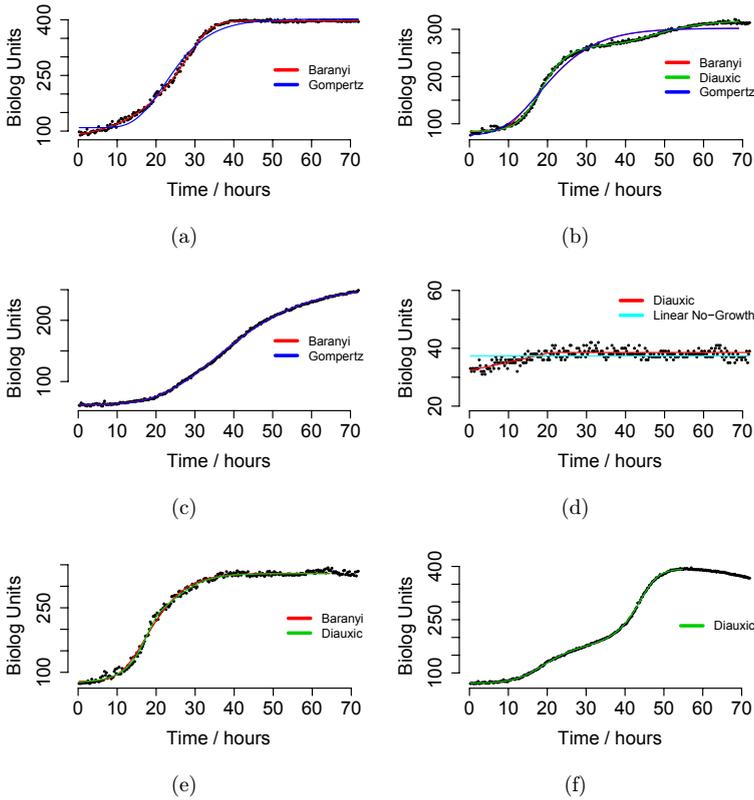


Fig. 3. Example data curves with model fits. (a) A typical data curve where the Baranyi model is best. The Baranyi model fits the data much better than the Gompertz model, that cannot fit either the shape of the lag phase or the shape of the transition to stationary phase. As a consequence, the Gompertz model would overestimate both the length of the lag phase and the maximal growth rate. (b) A typical data curve where diauxic model fits best. Neither the Baranyi nor the Gompertz model can capture the dynamics. (c) A curve with a preferred Gompertz model according to DIC. A fitted Baranyi-curve is also shown. (d) A curve showing no significant growth compared to a control well. Due to a slight increase in coloration this is fitted best equally by a diauxic and Baranyi model. Absence of growth is detected by comparing to the control well. (e) A borderline case that shows slight diauxic behavior and (f) Diauxic curve with high growth on  $s_2$ .

Baranyi and diauxic models receive much the best BIC scores (Table 1). Figure 3(c) shows one well with no growth, and a slight decrease in coloration, together with fits by the Baranyi and constant models.

Further evidence for the value of the Baranyi model and Baranyi lag term in the diauxic model in fitting this type of data can be seen from the relative wide spread of the estimated mean for the curvature parameter of the lag phase,  $\nu$  in both models, as well as the curvature parameter  $m$  in the Baranyi model (Fig. 4). Thus, we consider it likely that any model lacking such an extra parameter would fail to accommodate the range of curve shapes we have encountered.

Table 1. BIC scores for the relevant models in Fig. 3.

Strain	DIC Baranyi	DIC Diauxic	DIC Gompertz	DIC No-Growth
(a) P105-A12	388.5	398.0	2497.8	361073
(b) P106-E01	5457.6	<u>431.1</u>	5819.9	376824
(c) P107-G03	123.3	<u>123.1</u>	123.0	272838
(d) P106-H12	114.4	<u>114.2</u>	<u>120.8</u>	258
(e) P106-B11	548.1	<u>401.7</u>	555.5	135383
(f) P108-C10	3715.5	<u>142.9</u>	6502.7	491623

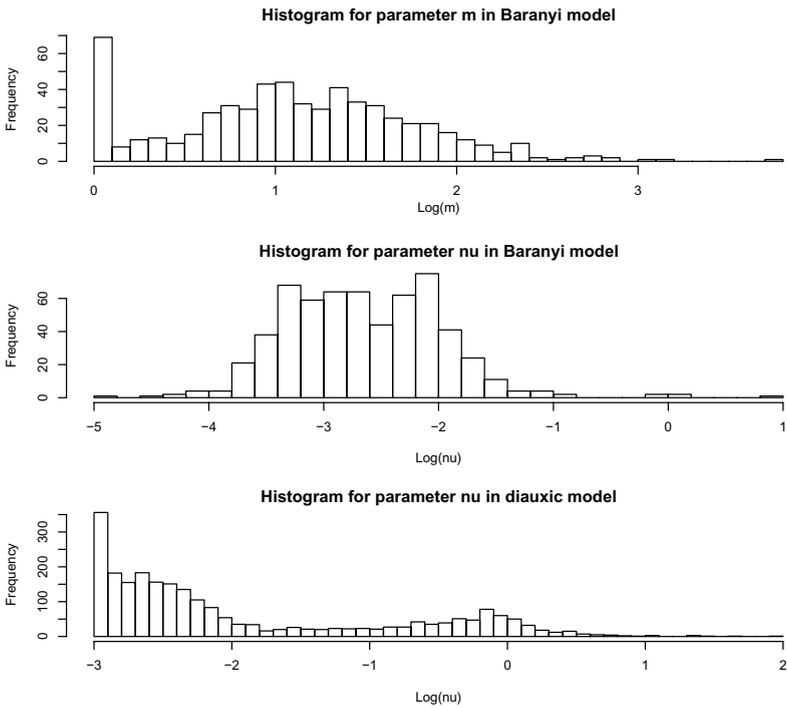


Fig. 4. Histograms for  $m$ ,  $\nu$  in the Baranyi and  $\nu$  in the diauxic model, for curves where these were the respective best-fit model and where significant growth was detected. The wide spread of these distributions is indicative of the importance of these parameters in fitting a range of data.

### 3.2. Identification of curve features beyond AUC

There are two main features of our approach: the first, as presented above, is that we are able to fit curves that are qualitatively diverse, namely nondiauxic and diauxic growth. The second is that we are able to estimate key relevant parameters from Biolog data, in particular a length of lag phase, maximal growth rate and maximal output. In many biotechnological applications, the identification of strains or conditions with minimal lag phase or maximal growth could be particularly important, as such strains or conditions could speed production and/or cut costs. Figures 5(a)

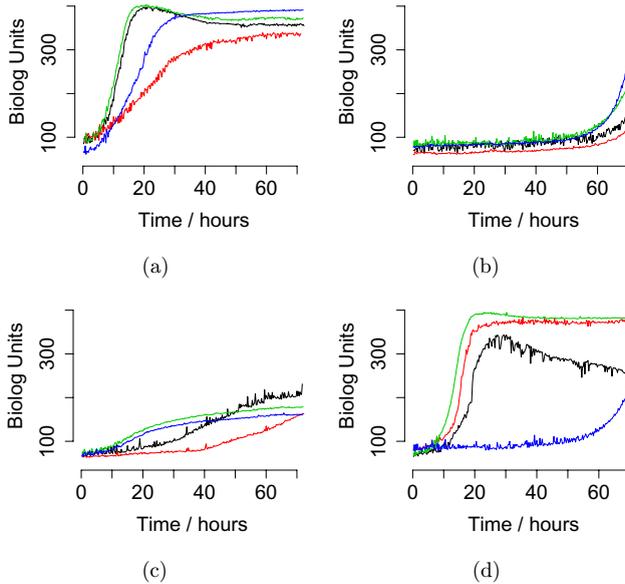


Fig. 5. Example data curves with potentially useful derived parameters. Some of the curves in the data set with (a) shortest lag phase, (b) longest lag phase, (c) lowest maximum growth rate (while still exhibiting significant growth) and (d) highest rate.

and 5(b) show the curves with the shortest and longest lag phases, respectively. These curves have quite different AUCs and maximal outputs so would be difficult to identify without a suitable modeling approach. Similarly, Figs. 5(c) and 5(d) show the curves with the fastest and slowest maximal growth. These too would be difficult to identify without a robust modeling approach.

Conversely, curves with very similar AUC can differ greatly in other parameters. Table 2 lists the mean derived parameters for the six different strains shown in the beginning of the paper, clearly showing that AUC alone is not sufficient to allow comparisons between qualitatively different data. The lag phases range from 12.3 h (P103-C01) to 38.9 h (P101-C01) and the maximal growth rates range from 11.7 per hour (P102-B12) to 35.0 per hour (P101-C01). Both of these model-derived parameters vary three-fold among the strains shown, demonstrating that our approach is considerably superior to using AUC.

We have found that a comparison to the 95% quantile of maximum coloration of a control well works reliably in identifying whether growth occurs.

### 3.3. The estimated parameters and model-independent measurements are reliable

The Bayesian approach has allowed us to use the posterior distributions to estimate standard deviations of the individual model parameters as well as the numerically derived model-independent growth measurements. These are generally low

Table 2. The AUC and mean estimated parameters for the four strains shown in Fig. 1.

Strain	AUC	L	$\mu_{\max}$	A
P101-D06	190.6	32.6	15.3	301.2
P102-E11	190.7	10.1	9.5	163.3
P103-C05	191.5	21.7	18.1	219.5
P104-C12	192.3	16.6	6.6	162.0

(typically  $< 1\%$  of the parameter values) and are elevated only in cases where some parameters are not estimable from the data, e.g. estimation of  $y_{\max}$  when fitting a Baranyi model to data that is cut off before stationary phase is reached (standard deviation as high as 10% of the parameter), but also sometimes  $s_1$  and  $s_2$  when fitting the diauxic model to data clearly depicting simple growth.

We have found that  $A$  (the maximum coloration change) very closely matches the respective model parameter(s), and that  $\mu_{\max}$  has given good comparability of results between wells even with different best-fit models. In cases of diauxic growth, particularly with small  $s_2$  and the bulk of growth on  $s_1$  (as is the case in the vast majority of diauxic cases observed) the Baranyi model would sometimes give an artificially low estimate of  $\mu_{\max}$  as it tries to compensate for the (in a sense slower) two-step approach to peak coloration by fitting an altogether slower curve. Figure 3(e) shows one such case together with fits by the Baranyi and diauxic models.

The lag time  $L$  in the vast majority of cases similarly gives good comparability between wells and models. In one atypical case 3(f), a diauxic growth curve features a higher growth rate on  $s_2$  than on  $s_1$ . By our definition of  $L$  as the time coordinate of the intersection of the steepest tangent with the flat line  $y \equiv y_0$ , the lag time in such cases is determined by the growth on  $s_2$  and could be significantly later than start of growth on  $s_1$ . Arguably this is not what we want, as it would not be consistent with a physiological definition of the lag phase, since the lag time identified includes the time grown on  $s_1$ . On the other hand, in some circumstances this output may be preferential, e.g. in an industrial research application a slow initial growth on  $s_1$  may be of considerably less interest than the time until peak growth rate. It would be possible in our framework to derive a specialized definition of  $L$  tailored to diauxic models, e.g. using the steepest tangent before the point where  $s_1$  is effectively depleted.

While our descriptive measurements  $L$ ,  $\mu_{\max}$  and  $A$  give us consistent definitions and readily comparable results for the vast majority of cases, the various model parameters we are also inferring allow us a more explanatory analysis of subsets of (or even individual) cases. For instance, within a set of diauxic curves, our estimates of  $s_1$  and  $s_2$  allow us to identify cases in which the bulk of growth is on the second substrate. Again, in industrial research applications this may be of particular interest. Similar analysis could be carried out e.g. on the curvature parameters of the Baranyi model, to identify outliers in lag phase behavior.

As a further test of reliability, we have compared parameter estimates from wells measuring the same conditions: each condition on these particular phenotype arrays

appears in triplicate. For the best fit data shown in Fig. 3, this provides three independent estimates for each of the three derived parameters for six models, a total of 18 comparisons. The parameter estimates were generally very consistent, with median percentage error of 5.8%. The worst case is for Fig. 3(d), where there is no growth, and the parameter estimates vary by approximately 30%; however, because there is no growth, this is not a problem. Full details of these comparisons are provided in Appendix D.

### 3.4. Performance

The MCMC methodology based on an Adaptive Metropolis algorithm performed well in combination with the models we used. We tested all the Markov chain outputs for the model fits that appear in all figures for convergence using the Heidelberger and Welch’s convergence test as implemented in the CODA package in R (<https://cran.r-project.org/web/packages/coda/index.html>). The results from this test showed that all the chains have passed the test, i.e. have converged to the stationary distribution. The outputs are not especially interesting (many tables of nonsignificant  $p$ -values) so we have placed one example in Appendix E and have not included the other outputs.

The adaptation to different target distributions in particular worked very well and our algorithm required no manual fine-tuning to explore different data with high efficiency. Figure 6 shows an initial segment of one AM chain we ran. It is clearly visible that the mixing of the chain improves rapidly after the initial discovery phase. Running our methodology implemented in C++ for a single plate (comprising 96 wells recorded every 15 min over 72 h) took around 50–60 min on a quad-core machine (Intel Xeon E3 1230v2).

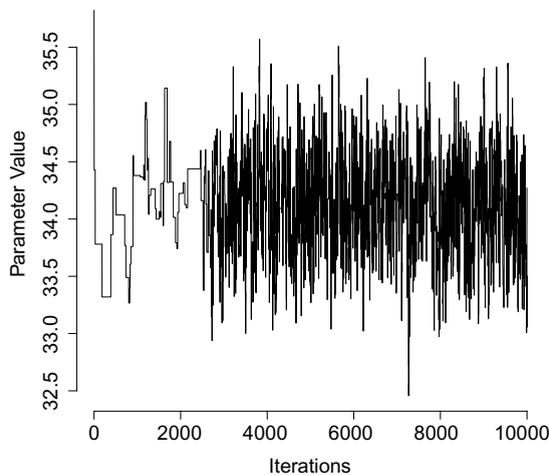


Fig. 6. Initial segment from the trace plot for  $\lambda$  for one AM chain. It is easy to see that the initial acceptance rate is far from optimal, but rapidly improves after the first 2500 iterations.

### 3.5. Comparison to previous approaches

Vaas *et al.* 2012 discuss one particular case (their Figs. 4(a) and 4(b)) where their model-fitting approach fares worse than their alternative spline-based method. It appears to us that the curve in question is simply diauxic in nature, and we expect that our approach would be able to fit this time course and identify the behavior as diauxic. The second example they discuss (Fig. 4(c) same curve) could equally likely be fitted by our methodology with an appropriate model. In either case Ref. 9 approach likely could not be extended easily to encompass such additional models, using a third-party software package to do the actual parameter estimation.

## 4. Conclusion

We have presented a Bayesian approach to estimating curve-parameter information from respiration data gathered in phenotype microarray experiments. We had aimed to extract meaningful information from complex respiration curves in a statistically sound way, and have shown that our approach succeeds in doing so.

Our solution has several key advantages. First, the Bayesian framework in which our approach operates affords us a great deal of flexibility in extracting information from respiration curves. As we are approximating the joint posterior distribution of all model parameters, we are free to perform further analysis on this and can for instance derive the distribution of arbitrary functions of our model parameters from them. Model choice criteria allow us to readily identify qualitatively different behavior in the data on top of quantitative measurements. This gives us a “best of two worlds” solution encompassing both unified model-independent descriptive measurements as well as more explanatory model-specific ones.

Second, the Baranyi model provides an excellent model to fit a wide range of observed growth curves. The only exception we encountered were curves exhibiting diauxie. In these cases, our diauxic model provided us with good fits in these cases as well. In particular, these two models outperform traditional growth models in many cases. Our modular implementation allows us to adapt our code to new types of experimental data with minimal effort. As PM machinery is being used in a wide range of different areas, this allows our solution to remain applicable as new types of data become available.

Third, the model-independent measurements  $L$ ,  $\mu_{\max}$  and  $A$  we derive from the fitted curves allow an immediate comparison of key values between different time series, even if qualitatively different models were used to fit the data. This allows us to match simpler (e.g. spline-based) approaches in their ability to perform quantitative comparisons between a wide range of growth behaviors. However in addition, our approach retains information from the richer model-dependent parameters. This gives us powerful tools to perform further analysis on specific subsets of data. In particular, we can use these to compare wells with qualitatively similar behavior in more detail, for instance utilization of different substrates in diauxic curves.

Last, our implementation is suited to high-throughput analysis of PM data. Compared to previous approaches to data analysis on PM experiments, we believe that our solution offers improvements in several areas. Even the relatively small number of models we have implemented so far allows our solution to successfully fit a wide range of data.

We believe that our extensible approach is uniquely suited to “keep up” with PM data, as PM machinery is being used for an increasingly wide field of different types of experiments. As PM machinery continues to find novel usage scenarios, the modular implementation we have chosen will allow our solution to remain flexible in accommodating data from these scenarios.

### **Program Code Availability**

We have made the program code available on GitHub at URL <https://github.com/dovstekellab/mcmc-pma.git> under license GNU GPLv3. We also include instructions on how to compile and run the code, as well as how to interpret the results files.

### **Acknowledgments**

We thank SABMiller for their support of this research project. Some preliminary work on fitting growth models to PM data using an MCMC approach was also carried out by Ewan Johnstone on a BBSRC-funded summer studentship at the University of Nottingham. MG is currently recipient of a DOC Fellowship of the Austrian Academy of Sciences.

### **Appendix A. Prior Distributions**

We used uninformative or uniform priors as follows:

- (1) For models that feature a lag parameter  $\lambda$ , we assumed  $0 \leq \lambda \leq \dim(x)$ .
- (2) For all parameters where we were not dealing with their logarithm anyways, we assumed that they are positive.
- (3) Where applicable we assumed  $y_{\max} > y_0$ .
- (4) In the full Baranyi model, we required  $r > \nu$ ,  $m > 1$ , as otherwise these parameters were usually inestimable due to high correlation. Similarly in the Baranyi models where only one of  $\nu$ ,  $m$  are fixed.
- (5) For all other parameters, we used improper uniform priors without any bounds.

### **Appendix B. Estimation of Initial Parameters**

The heuristics we used to estimate initial parameter vectors are as follows:

- (1) For initial cell concentration, we took the average of the first 10 data points. Similarly for the maximum concentration, we took the average of the maximum of any 10-point interval in the time series.

- (2) For the lag parameter, we fit lines to every 20-point interval of data points, and intersected the maximum-slope line with  $y = y_0$ , with  $y_0$  derived as in the previous line.
- (3) For the growth rate  $r$ , we used a simple search heuristic to guess an initial value that we have found empirically to be effective. We start from a sufficiently large interval of possible values  $[r_{\min}, r_{\max}]$ , and divided this into 10 equal parts  $r_0 = r_{\min}, r_1, \dots, r_{11} = r_{\max}$ . We then compared the maximum slope of the modeled data with  $r = r_1, \dots, r_{10}$  with the maximum slope as in the previous line. For  $r_i$  giving the smallest difference in maximum slope, we recurse by setting  $r_{\min} = r_{i-1}, r_{\max} = r_{i+1}$ . We repeat this 100 times.
- (4) For the remaining parameters we used the same search algorithm, except we compared the sum of squares of differences between modeled and observed data instead of maximum slope.

Specific to each of the models we did the following: We used  $[0.01, 1]$ ,  $[2, 10]$ ,  $[0.3, 1]$  as initial intervals for  $r$ ,  $m$ ,  $\nu$  in this search algorithm. We first called this for  $r$ , setting both remaining parameters to 1, then similarly for the remaining ones. For the diauxic models, we proceed similarly, except we additionally needed to find  $s_1$  and  $s_2$ , respectively the division of the total growth between the two. To do so, we found the maximum slope on a smoothed curve as in our estimation for the lag parameter, and then the minimum slope between that point and when the smoothed curve first comes within 12 standard deviations of the maximum. (In other words, we looked for the characteristic intermittent slowing-down of growth.) We then subtracted  $y_0$  from the value at this point, and took this as the initial guess for  $s_1$ . For the remaining parameters we proceeded similarly as in the Baranyi model, using  $[10^{-8}, 10^{-4}]$ ,  $[0.05, 1]$ ,  $[0.0001, 11]$  as initial intervals for  $r_1$ ,  $\nu$ ,  $k_1$  in the search algorithm. We always use 0.0003 for  $r_2$ , as performing our search heuristic for this parameter did not improve results.

### Appendix C. Estimation of Measurement Error from Data

We took a smoothed curve (taking the average of every nine adjacent measurements; without taking logarithms) as reference to numerically estimate the variance. We however always assumed a minimum variance of 5 Biolog units. That is, we defined

$$\sigma^2 = \min \left\{ 5, \frac{1}{\dim(x)} \sum_{j=4}^{\dim(x)-5} x_j - \frac{1}{9} (x_{j-4} + x_{j-3} + \dots + x_{j+4}) \right\}.$$

### Appendix D. Consistency Analysis

Consistency of parameter estimates was tested for the six best-fit models shown in Fig. 3. These particular phenotype microarrays have triplicate wells for each condition,

Table D.1. Parameter estimate consistency.

Wells	Parameter	Value1	Value2	Value3	Mean	St. Dev.	% Error	Note
P105-A10-A11-A12	lag	15.37	14.37	14.58	14.77	0.530	3.58%	
	rate	17.60	17.21	15.79	16.87	0.951	5.64%	
	max	325.3	321.2	306.8	317.8	9.69	3.05%	
P106-E01-E02-E03	lag	11.79	11.31	11.80	11.63	0.284	2.44%	
	rate	14.03	11.78	12.00	12.60	1.24	9.83%	
	max	232.4	210.4	201.4	214.7	16.0	7.44%	
P107-G01-G02-G03	lag	20.38	24.57	20.36	21.77	2.43	11.1%	
	rate	8.06	7.06	5.17	6.76	1.47	21.7%	
	max	186.4	183.1	183.6	184.3	1.76	0.96%	
P106-H10-H11-H12	lag	3.64	5.64	3.04	4.11	1.36	33.1%	No growth case
	rate	0.634	0.452	0.436	0.507	0.110	21.7%	No growth case
	max	8.90	5.05	5.53	6.49	2.10	32.3%	No growth case
P106-B10-B11-B12	lag	11.83	11.80	11.58	11.74	0.135	1.15%	
	rate	19.33	19.41	21.46	20.07	1.21	6.02%	
	max	256.0	253.2	265.6	258.3	6.48	2.51%	
P108-C10-C11-C12	lag	31.30	30.10	28.59	30.00	1.36	4.52%	
	rate	18.74	15.79	15.07	16.53	1.95	11.8%	
	max	320.6	310.9	315.8	315.8	4.86	1.54%	

so the derived model parameters shown were compared with those for the two other triplicate wells. The full output is:

### Appendix E. Convergence Tests

Convergence for all of the Markov chains for model fits used in the figures was carried out using the Heidelberger and Welch’s convergence test as implemented in the CODA package in R (<https://cran.r-project.org/web/packages/coda/index.html>). All of the chains passed the test, i.e. demonstrating convergence to a stationary distribution. The output for each chain is in the form of a table with a  $p$ -value for each parameter. An example from one chain that is typical of all output is:

Each row, labeled V1 to V8, represents one variable. The null hypothesis is that the chain is from a stationary distribution and so it can be seen that the null

Table E.1. Example convergence test output from CODA.

Variable	Stationarity test	Start iteration	$p$ -value
V1	passed	1	0.357
V2	passed	1	0.547
V3	passed	1	0.560
V4	passed	1	0.297
V5	passed	1	0.602
V6	passed	1	0.853
V7	passed	1	0.218
V8	passed	1	0.173

hypothesis has been accepted for all variables in the chain. We obtained similar results for all of the Markov chains tests (data not shown).

## References

1. Andrieu C, Atchadé Y, On the efficiency of adaptive mcmc algorithms, *Proc 1st Int Conf Performance Evaluation Methodologies and Tools*, ACM, p. 43, 2006.
2. Andrieu C, Thoms J, A tutorial on adaptive mcmc, *Stat Comput* **18**(4):343–373, 2008.
3. Atchadé Y, Fort G, Limit theorems for some adaptive mcmc algorithms with subgeometric kernels, *Bernoulli* **16**(1):116–154, 2010.
4. Bai Y, Roberts G, Rosenthal J, On the containment condition for adaptive Markov chain Monte Carlo algorithms, *Technical Report*, University of Warwick. Centre for Research in Statistical Methodology, 2009.
5. Baranyi J, Simple is good as long as it is enough, *Food Microbiol* **14**(4):391–394, 1997.
6. Baranyi J, Roberts T, A dynamic approach to predicting bacterial growth in food, *Int J Food Microbiol* **23**(3):277–294, 1994.
7. Baranyi J, Roberts T, Mathematics of predictive food microbiology, *Int J Food Microbiol* **26**(2):199–218, 1995.
8. Baranyi J, Roberts T, McClure P, A non-autonomous differential equation to model bacterial growth, *Food Microbiol* **10**(1):43–59, 1993.
9. Bochner B, New technologies to assess genotype–phenotype relationships, *Nat Rev Genet* **4**(4):309–314, 2003.
10. Bochner B, Global phenotypic characterization of bacteria, *FEMS Microbiol Rev* **33**(1):191–205, 2009.
11. Bochner B, Gadzinski P, Panomitros E, Phenotype microarrays for high-throughput phenotypic testing and assay of gene function, *Genome Res* **11**(7):1246–1255, 2001.
12. DeNittis M, Querol A, Zanoni B, Minati JL, Ambrosoli R, Possible use of biolog methodology for monitoring yeast presence in alcoholic fermentation for wine-making, *J Appl Microbiol* **108**:1199–1206, 2009.
13. DeNittis M, Zanoni B, Minati JL, Gorra R, Ambrosoli R, Modelling biolog profiles’ evolution for yeast growth monitoring in alcoholic fermentation, *Lett Appl Microbiol* **52**:96–103, 2010.
14. Eddy SR, What is bayesian statistics?, *Nature Biotechnol* **22**(9):1177–1178, 2004.
15. Gamerman D, Lopes H, *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, Chapman & Hall/CRC, UK, 2006.
16. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB, *Bayesian Data Analysis*, CRC press, UK, 2013.
17. Gilks W, Richardson S, Spiegelhalter D, *Markov Chain Monte Carlo in Practice*, Chapman & Hall/CRC, UK, 1996.
18. Haario H, Saksman E, Tamminen J, An adaptive metropolis algorithm, *Bernoulli* **7**(2):223–242, 2001.
19. Hastings WK, Monte carlo sampling methods using markov chains and their applications, *Biometrika* **57**(1):97–109, 1970.
20. Jacobsen J, Joyner D, Borglin S, Hazen T, Arkin A, Bethel E, Visualization of growth curve data from phenotype microarray experiments, *11th Int Conf Information Visualization, 2007. IV’07, IEEE*, pp. 535–544, 2007.
21. Kahm M, Hasenbrink G, Lichtenberg-Fraté H, Ludwig J, Kschischo M, grofit: Fitting biological growth curves with r, *J Stat Softw* **33**(7):1–21, 2010.

22. Kompala D, Ramkrishna D, Jansen N, Tsao G, Investigation of bacterial growth on mixed substrates: Experimental evaluation of cybernetic models, *Biotechnol Bioeng* **28**(7):1044–1055, 1986.
23. Kompala D, Ramkrishna D, Tsao G, Cybernetic modeling of microbial growth on multiple substrates, *Biotechnol Bioeng* **26**(11):1272–1281, 1984.
24. Lopez S, Prieto M, Dijkstra J, Dhanoa M, France J, Statistical evaluation of mathematical models for microbial growth, *Int J food Microbiol* **96**(3):289–300, 2004.
25. Peleg M, Corradini MG, Microbial growth curves: What the models tell us and what they cannot, *Crit Rev Food Sci Nutr* **51**(10):917–945, 2011.
26. Richards F, A flexible growth function for empirical use, *J Exp Bot* **10**(2):290–301, 1959.
27. Schwarz G, Estimating the dimension of a model, *Ann Stati* **6**(2):461–464, 1978.
28. Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A, Bayesian measures of model complexity and fit, *J R Stat Soc, B Stati Methodol* **64**(4):583–639, 2002.
29. Torres NP, Lee AY, Giaever G, Nislow C, Brown GW, A high-throughput yeast assay identifies synergistic drug combinations, *Assay Drug Dev Technol* **11**(5):299–307, 2013.
30. Vaas L, Sikorski J, Michael V, Göker M, Klenk H, Visualization and curve-parameter estimation strategies for efficient exploration of phenotype microarray kinetics, *PLoS one* **7**(4):e34846, 2012.
31. Vehkala M, Shubin M, Connor TR, Thomson NR, Corander J, Novel R, Pipeline for Analyzing Biolog Phenotypic Microarray Data, *PLoS one* **10**(3):e0118392, 2015.
32. Zhou L, Lei X, Bochner B, Wanner B, Phenotype microarray analysis of escherichia coli k-12 mutants with deletions of all two-component systems, *J Bacteriol* **185**(16):4956–4972, 2003.
33. Zwietering M, Jongenburger I, Rombouts F, Van't Riet K, Modeling of the bacterial growth curve, *Appl Environ Microbiol* **56**(6):1875–1881, 1990.



**Matthias Gerstgrasser** received a Mag.phil. degree from the University of Vienna in 2010, a MAST from Cambridge University in 2011 and a MRes from the University of Nottingham in 2012. During 2012–2013 he was a member of the European Space Agency's Advanced Concepts Team in the Netherlands. Since 2013 he has been a PhD student in Computer Science at Oxford University. His current research focuses on CS Theory, and in particular Algorithmic Game Theory.

**Sarah Nicholls** received her BSc. in Applied Biology from the University of Wales Institute of Science and Technology, (UWIST), in Cardiff (1988). She completed a PhD in molecular biology whilst working at that Center for Applied Microbiology and Research (CAMR), Porton Down (1988–1994). After a long career break, she began post-doctoral research in Food Sciences at the University of Nottingham (2006) and from 2008 has worked in Brewing Science. Since 2010 her research has focused on investigation of the properties of the many strains of commercial brewing yeast held by the brewing giant SABMiller.



**Michael Stout** has a background in Evolutionary Computation, Functional Programming, Information Visualization and High Performance Computation in particular using GPUs. He is currently a Post Doctoral Research Fellow at the DeepSeq Facility in the Faculty of Medicine at University of Nottingham, engaged on a BBSRC funded project to develop tools and techniques for high throughput, real time processing and control of nanopore DNA sequencing devices. Prior to this he worked on Bayesian model fitting and comparison for Biolog and related datasets, funded by BBSRC.



**Katherine Smart** completed a BSc (Hons) in Biological Sciences at Nottingham University and was awarded the Rainbow Research Scholarship to complete a PhD in Brewing Yeast and Fermentation at Bass Brewers, Burton-on-Trent. This was where her passion for brewing science began. She has held fellowships and academic positions at the University of Cambridge, Oxford Brookes University and most recently at the University of Nottingham where she was the SABMiller Professor of Brewing Science. Katherine cites as her achievements the founding of the National Brewing Library at Oxford Brookes University and the Bioenergy and Brewing Science Centre at the University of Nottingham. During her academic career Katherine received several awards for her research including the Institute of Brewing and Distilling Cambridge Prize (1999), the prestigious Royal Society Industrial Fellowship (2001–2003), an Enterprise Fellowship (2002) and the Save British Science Award at the Houses of Parliament in the UK (2003). In 2012, she joined SABMiller as Group Chief Brewer. She also holds a Professorial role at the University of Nottingham. Katherine is a Fellow of the Institute of Brewing and Distilling, the Royal Society of the Arts, Manufacturing and Commerce and the Society of Biology. She has served as Council and Board member for the IBD, was International Director for the ASBC and is a long standing member of the MBAA.



**Chris Powell** holds a PhD on yeast cellular ageing and fermentation performance from Oxford Brookes University, UK (2001). After a post-doctoral research position at the same institute investigating rapid methods for detection of contaminants in beer (2001–2004) he worked in the research department for Lallemand Inc, Montreal, Canada (2004–2010), rising to occupy the positions of Senior Scientist in Brewing and Project Manager in Genetic Identification. In 2010 Chris moved to the University of Nottingham where he is a lecturer in yeast and fermentation. His core subject areas include yeast physiology and fermentation biotechnology, particularly related to the beverage and sustainable bioenergy sectors. Chris has held a position on the American Society of Brewing Chemists (ASBC) Board of Directors since 2010, currently acting as Vice President.

**Theodore Kypraios** is an Associate Professor in Statistics in the School of Mathematical Sciences, at the University of Nottingham. His research is mostly concerned with the development of novel methodology for efficient Bayesian statistical inference and model selection for complex high-dimensional datasets. The area that he has worked the most is on formulating stochastic epidemic models and fitting them to disease outbreak data of infectious diseases using Monte Carlo methods. During recent years he has expanded his research interests outside epidemic modelling and has developed and applied statistical methodology to a wide range of inter-disciplinary research projects including systems biology, phenotype microarray data neuroimaging, ecological modelling, financial modelling, internet traffic modelling, quantum statistics and RNA-Sequence expression data.



**Dov Stekel** is Associate Professor of Integrative Systems Biology in the School of Biosciences, University of Nottingham. He studied mathematics at Trinity College, Cambridge and for a D.Phil. at the Department of Zoology, University of Oxford. He worked briefly at the Gray Laboratory, before taking a position in the Bioinformatics group at Glaxo Wellcome Research and Development (as it was), and then as Head of Bioinformatics at Oxford Gene Technology. He left OGT after receiving a commission from Cambridge University Press to write *Microarray Bioinformatics*. In 2004 he took the position as Lecturer in Bioinformatics at the University of Birmingham and in 2009 moved to the University of Nottingham to his current post.