Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **474,** 2094–2108 (2018) doi:10.1093/mnras/stx2109
Advance Access publication 2017 September 22

Inferring probabilistic stellar rotation periods using Gaussian processes

Ruth Angus,[1]★ Timothy Morton,[2]★ Suzanne Aigrain,[3] Daniel Foreman-Mackey[4] and Vinesh Rajpaul[3]

[1]*Simons Fellow, Department of Astronomy, Columbia University, NY 10027, USA*
[2]*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*
[3]*Subdepartment of Astrophysics, University of Oxford, Oxford, OX1 3RH, UK*
[4]*Sagan Fellow, Department of Astronomy, University of Washington, Seattle, WA 98195, USA*

Accepted 2017 August 11. Received 2017 August 11; in original form 2017 March 14

ABSTRACT

Variability in the light curves of spotted, rotating stars is often non-sinusoidal and quasi-periodic – spots move on the stellar surface and have finite lifetimes, causing stellar flux variations to slowly shift in phase. A strictly periodic sinusoid therefore cannot accurately model a rotationally modulated stellar light curve. Physical models of stellar surfaces have many drawbacks preventing effective inference, such as highly degenerate or high-dimensional parameter spaces. In this work, we test an appropriate *effective* model: a Gaussian Process with a quasi-periodic covariance kernel function. This highly flexible model allows sampling of the posterior probability density function of the periodic parameter, marginalizing over the other kernel hyperparameters using a Markov Chain Monte Carlo approach. To test the effectiveness of this method, we infer rotation periods from 333 simulated stellar light curves, demonstrating that the Gaussian process method produces periods that are more accurate than both a sine-fitting periodogram and an autocorrelation function method. We also demonstrate that it works well on real data, by inferring rotation periods for 275 *Kepler* stars with previously measured periods. We provide a table of rotation periods for these and many more, altogether 1102 *Kepler* objects of interest, and their posterior probability density function samples. Because this method delivers posterior probability density functions, it will enable hierarchical studies involving stellar rotation, particularly those involving population modelling, such as inferring stellar ages, obliquities in exoplanet systems, or characterizing star–planet interactions. The code used to implement this method is available online.

Key words: methods: data analysis – methods: statistical – techniques: photometric – stars: rotation – stars: solar-type – starspots.

1 INTRODUCTION

The disc-integrated flux of a spotted, rotating star often varies in a non-sinusoidal and quasi-periodic (QP) manner, due to active regions on its surface that rotate in and out of view. Complicated surface spot patterns produce non-sinusoidal variations, and the finite lifetimes of these active regions and differential rotation on the stellar surface produce quasi-periodicity (Dumusque et al. 2011). A strictly periodic sinusoid is therefore not necessarily a good model choice for these time series. A physically realistic model of the stellar surface would perfectly capture the complexity of shapes within stellar light curves as well as the QP nature, allowing for extremely precise probabilistic period recovery when conditioned on the data. However, such physical models require many free parameters in order to accurately represent a stellar surface, and some of these parameters are extremely degenerate (e.g. Russell 1906; Jeffers & Keller 2009; Kipping 2012). In addition to global stellar parameters such as inclination and rotation period, each spot or active region should have (at minimum) a longitude, latitude, size, temperature and lifetime. Considering that stars may have hundreds of spots, the number of free parameters in such a model quickly becomes unwieldy, especially to explore its posterior probability density function (PDF). Simplified spot models, such as the one described in Lanza, Das Chagas & De Medeiros (2014) where only two spots are modelled, have produced successful results; however, such relatively inflexible models sacrifice precision.

★ E-mail: ruthangus@gmail.com (RA); tim.morton@gmail.com (TM)

© 2017 The Author(s)
Published by Oxford University Press on behalf of the Royal Astronomical Society

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

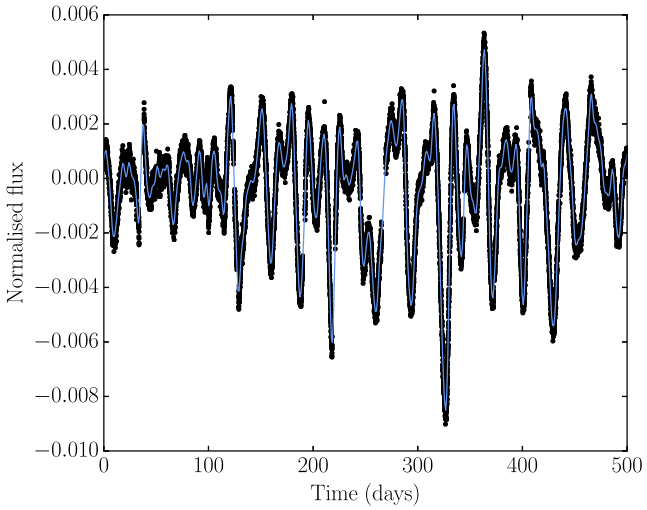

Figure 1. Light curve of KIC 5809890, an active star with a rotation period of $\sim$ 30.5 d. The solid line shows a fit to the data using a GP model with a QP covariance kernel function.

Standard non-inference-based methods to measure rotation periods include detecting peaks in a Lomb–Scargle (LS; Lomb 1976; Scargle 1982) periodogram (e.g. Reinhold, Reiners & Basri 2013), autocorrelation functions (ACFs; e.g. McQuillan, Mazeh & Aigrain 2013b) and wavelet transforms (e.g. García et al. 2014). The precisions of the LS periodogram and wavelet methods are limited by the suitability of the model choice: a sinusoid for the LS periodogram, and a choice of mother wavelet, assumed to describe the data over a range of transpositions and scales (see e.g. Carter & Winn 2010), for the wavelet method. In contrast, since it does not rely on a generative model, the ACF method is much better suited to signals that are non-sinusoidal. In fact, as long as the signal is approximately periodic the ACF will display a peak at the rotation period, no matter its shape. A drawback of the ACF method, however, is that it requires data to be evenly spaced. Most ground-based measurements, such as the future Large Synoptic Survey Telescope (*LSST*), do not nearly satisfy this criterion, and even light curves from the *Kepler* space telescope, which has provided the richest stellar rotation data to date, can only be approximated as uniformly sampled. An ACF is also an operation performed on the data rather than a generative model of the data, and so is not inherently probabilistic. This means that the effects of the observational uncertainties cannot be formally propagated to constraints on the rotation period.

In this work, we test an *effective* model for rotationally modulated stellar light curves that captures the salient behaviour but is not physically motivated – although some parameters may indeed be *interpreted* as physical ones. An ideal effective model should have a small number of non-degenerate parameters and be flexible enough to perfectly capture non-sinusoidal and QP behaviour. A Gaussian process (GP) model fulfills these requirements. We thus use a GP as the generative model at the core of a method to probabilistically infer accurate and precise stellar rotation periods. This enables us to estimate the posterior PDF of the rotation period, thereby producing a justified estimate of its uncertainty. Using a quasi-periodic Gaussian process (QP-GP) model to infer a rotation period is not a new idea. For example, we used this method in Vanderburg et al. (2015) to measure the rotation period of an exoplanet host. Previously to that, Haywood et al. (2014) used the QP-GP model to disentangle the correlated noise produced by stellar activity and rotation from the radial velocity signature of an exoplanet. In addition, Littlefair, Burningham & Helling (2017) use a QP-GP to establish the quasi-periodicity of variability in the light curves of brown dwarfs. What we present here is a test of the GP method and a comparison with alternative rotation period measurement methods. We also provide code allowing others to easily apply this technique themselves.

GPs are commonly used in the machine learning community and increasingly in other scientific fields such as biology, geophysics and cosmology. More recently, GPs have been used in the stellar and exoplanet fields within astronomy to capture stellar variability or instrumental systematics (see e.g. Gibson et al. 2012; Dawson et al. 2014; Haywood et al. 2014; Barclay et al. 2015; Czekala et al. 2015; Evans et al. 2015; Haywood 2015; Rajpaul et al. 2015; Vanderburg et al. 2015; Aigrain, Parviainen & Pope 2016; Rajpaul, Aigrain & Roberts 2016; Littlefair et al. 2017). They are useful in regression problems involving any stochastic process, specifically when the probability distribution for the process is a multivariate Gaussian. If the probability of obtaining a data set is a Gaussian in N dimensions, where N is the number of data points, a GP can describe that data set. An in-depth introduction to GPs is provided in Rasmussen & Williams (2005).

GP models parametrize the covariance between data points by means of a kernel function. As a qualitative demonstration, we present the time series in Fig. 1: the *Kepler* light curve of KIC 5809890. This is a relatively active star that rotates once every $\sim$30.5 days, with stochastic variability typical of *Kepler* FGK stars. Clearly, data points in this light curve are correlated. Points close together in time are tightly correlated, and points more widely separated are loosely correlated. A GP models this variation in correlation as a function of the separation between data points; that is, it models the *covariance structure* rather than the data directly. This lends GPs their flexibility – they can model any time series with a similar covariance structure. In addition, a very simple function can usually capture the covariance structure of a light curve, whereas modelling the time series itself might require much more complexity. Fig. 1 demonstrates how a GP model fits the light curve of KIC 5809890.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

A range of covariance models, or kernel functions, can describe stellar variability. For example, the most commonly used kernel function, the 'squared exponential' (SE), defined as follows, could adequately fit the KIC 5809890 light curve:

$$k_{i,j} = A \exp\left(-\frac{(x_i - x_j)^2}{2l^2}\right). \quad (1)$$

Here, $A > 0$ is the amplitude of covariance, l is the lengthscale of exponential decay, and $x_i - x_j$ is the separation between data points. The SE kernel function has the advantage of being very simple, with just two parameters, A and l. If l is large, two data points far apart in x will be tightly correlated, and if small they will be loosely correlated. Another property of the SE kernel function is that it produces functions that are infinitely differentiable, making it possible to model a data set and its derivatives simultaneously. However, the SE kernel function does not well describe the covariance in stellar light curves, nor is it *useful* for the problem of rotation period inference because it does not capture periodic behaviour. Inferring rotation periods thus requires a periodic kernel function. For this reason, we use the 'QP' kernel. Rasmussen & Williams (2005) model QP variability in CO_2 concentration on the summit of the Mauna Loa volcano in Hawaii (data from Keeling & Whorf 2004) using a kernel which is the product of a periodic and a SE kernel: the QP kernel. This kernel is defined as

$$k_{i,j} = A \exp\left[-\frac{(x_i - x_j)^2}{2l^2} - \Gamma^2 \sin^2\left(\frac{\pi(x_i - x_j)}{P}\right)\right] + \sigma^2 \delta_{ij}. \quad (2)$$

It is the product of the SE kernel function, which describes the overall covariance decay, and an exponentiated, squared, sinusoidal kernel function that describes the periodic covariance structure. P can be interpreted as the rotation period of the star, and Γ controls the amplitude of the $\sin^2$ term. If Γ is very large, only points almost exactly one period away are tightly correlated and points that are slightly more or less than one period away are very loosely correlated. If Γ is small, points separated by one period are tightly correlated, and points separated by slightly more or less are still highly correlated, although less so. In other words, large values of Γ lead to periodic variations with increasingly complex harmonic content. This kernel function allows two data points that are separated in time by one rotation period to be tightly correlated, while also allowing points separated by half a period to be weakly correlated. The additional parameter σ captures white noise by adding a term to the diagonal of the covariance matrix. This can be interpreted to represent underestimation of observational uncertainties – if the uncertainties reported on the data are too small, it will be non-zero – or it can capture any remaining 'jitter', or residuals not captured by the effective GP model. We use this QP kernel function (equation 2) to produce the GP model that fits the *Kepler* light curve in Fig. 1.

There are many ways to construct a QP kernel function, involving a range of choices for both the periodic and aperiodic components of the model. The kernel function presented above reproduces the behaviour of stellar light curves, but other model choices can do so as well. We do not attempt to test any other models in this paper, noting only that this kernel function provides an adequate fit to the data. We leave formal comparisons with other kernel function choices to a future publication.

To infer a stellar rotation period P from a light curve, we fit this QP-kernel GP model to the data. As with any model-fitting exercise, the likelihood of the model could be maximized to find the maximum-likelihood value for P. In this study, however, we explore the full posterior PDFs using a Markov Chain Monte Carlo (MCMC) procedure. While this approach comes at a computational cost, such posterior exploration importantly provides a justified uncertainty estimate.

Performing inference with GPs is computationally expensive. Is it worth spending the extra computation power to get slightly more accurate, probabilistic rotation periods for a large number of stars? A growing number of astronomers are interested in hierarchical probabilistic modelling, particularly in the stellar and exoplanet communities. These individuals are mindful of carefully inferring probabilistic parameters wherever possible. Probabilistic inference allows one to marginalize over nuisance parameters while formally taking any uncertainty into account when measuring the quantity of interest – the star's rotation period, in this case. This becomes especially important as the signal-to-noise ratio decreases and the rotation period is no longer tightly constrained, and as the sampling becomes uneven. Throughout the astronomical literature, hierarchical Bayesian inference has been demonstrated to be a robust method of population inference (e.g. Hogg, Myers & Bovy 2010; Foreman-Mackey, Hogg & Morton 2014b; Rogers 2015; Wolfgang, Rogers & Ford 2016). Probabilistic measurements with full, rigorous propagation of uncertainties are a crucial ingredient for any hierarchical inference where the target is the population level distribution. Inferring the ages of stars in the galaxy via gyrochronology (age–rotation relations) (e.g. Skumanich 1972; Kawaler 1989; Barnes 2003, 2007) is an example of a hierarchical inference problem using stellar rotation. Hierarchical inference allows you to marginalize over individual rotation periods and infer age distribution parameters directly from the light curves. Probabilistic rotation period inference is highly relevant to a large number of galactic archaeology and exoplanet population studies that require rotation period (or age) distributions. It is important to note, however, that gyrochronology may have limited applicability for galactic field stars: growing evidence points to weakened magnetic braking at late ages (Angus et al. 2015; van Saders et al. 2016).

This paper is laid out as follows. The GP method is described in Section 2. Its performance is demonstrated and compared with literature methods in Section 3. In Section 4, we apply the GP method to real *Kepler* data, and the results are discussed in Section 5.

2 GP ROTATION PERIOD INFERENCE

In order to recover a stellar rotation period from a light curve using a QP-GP, we sample the following posterior PDF:

$$p(\boldsymbol{\theta} \mid y) \propto \mathcal{L}(\boldsymbol{\theta}) p(\boldsymbol{\theta}), \quad (3)$$

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

where y are the light-curve flux data, $\boldsymbol{\theta}$ are the hyperparameters of the kernel described in equation (2), $\mathcal{L}$ is the QP-GP likelihood function and $p(\boldsymbol{\theta})$ is the prior on the hyperparameters. Sampling this posterior presents several challenges:

(i) The likelihood evaluation is computationally expensive.
(ii) The GP model is very flexible, sometimes at the expense of reliable recovery of the period parameter.
(iii) The posterior may often be multimodal.

This section discusses how we address these challenges through the implementation details of the likelihood (Section 2.1), priors (Section 2.2) and sampling method (Section 2.3).

2.1 Likelihood

The GP likelihood is similar to the simple Gaussian likelihood function, where the uncertainties are Gaussian and uncorrelated. The latter can be written as

$$\ln \mathcal{L} = -\frac{1}{2}\sum_{n=1}^{N}\left[\frac{(y_n-\mu)^2}{\sigma_n^2} + \ln(2\pi\sigma_n^2)\right], \tag{4}$$

where y_n are the data, μ is the mean model and σ_n are the Gaussian uncertainties on the data. The equivalent equation in matrix notation is

$$\ln \mathcal{L} = -\frac{1}{2}\boldsymbol{r}^T\mathbf{C}^{-1}\boldsymbol{r} - \frac{1}{2}\ln|\mathbf{C}| + \frac{N}{2}\log 2\pi, \tag{5}$$

where $\boldsymbol{r}$ is the vector of residuals and $\mathbf{C}$ is the covariance matrix,

$$\mathbf{C} = \begin{pmatrix} \sigma_1^2 & \sigma_{2,1} & \cdots & \sigma_{N,1} \\ \sigma_{1,2} & \sigma_2^2 & \cdots & \sigma_{N,2} \\ & & \vdots & \\ \sigma_{1,N} & \sigma_{2,N} & \cdots & \sigma_N^2 \end{pmatrix} \tag{6}$$

In the case where the uncertainties are uncorrelated, the noise is 'white' (an assumption frequently made by astronomers, sometimes justified) and the off-diagonal elements of the covariance matrix are zero. However, in the case where there is evidence for correlated 'noise',[1] as in the case of *Kepler* light curves, those off-diagonal elements are non-zero. With GP regression, a covariance matrix generated by the kernel function $\mathbf{K}$ replaces $\mathbf{C}$ in the above equation (for our purposes, the QP kernel of equation 2 generates $\mathbf{K}$).

Evaluating this likelihood for a large number of points can be computationally expensive. For example, evaluating $\mathcal{L}$ for an entire *Kepler* light curve (∼40 000 points) takes about ∼5 s – too slow to perform inference on large numbers of light curves.[2] The matrix operations necessary to evaluate the GP likelihood scale as $N\ln(N)^2$, where N is the number of data points in the light curve, using the fast matrix solver HODLR (Ambikasaran et al. 2015), implemented in the `george` (Foreman-Mackey et al. 2014a) PYTHON package.

We accelerate the likelihood calculation using two complementary strategies: subsampling the data and splitting the light curve into independent sections. To subsample *Kepler* data, for example, we randomly select 1/30th of the points in the full light curve (an average of ∼1.5 points per day). This decreases the likelihood evaluation time by a factor of about 50, down to about 100 ms. We then split the light curve into equal-sized chunks containing approximately 300 points per section (corresponding to about 200 d) and evaluate the log-likelihood as the sum of the log-likelihoods of the individual sections (all using the same parameters $\boldsymbol{\theta}$). This reduces computation time because the section-based likelihood evaluation scales as $mn\ln(n)^2$, where n is the number of data points per section and m is the number of sections. This method further reduces computation time for a typical light curve (subsampled by a factor of 30) by about a factor of 2, to about 50 ms.

2.2 Priors

2.2.1 Non-period hyperparameters

The flexibility of this GP model allows for posterior multimodality and 'over-fitting'-like behaviour. For example, if l is small, then non-periodic factor in the covariance kernel may dominate, allowing for a good fit to the data without requiring any periodic covariance structure – even if clear periodic structure is present. Additionally, for large values of Γ the GP model becomes extremely flexible and can fit the data without varying the period. Managing this flexibility to reliably retrieve the period parameter requires imposing informative priors on the non-period GP parameters. In particular, we find it necessary to avoid large values of A and Γ, and small values of l; though the exact details of what works well may differ among data sets.

[1] In our case, the 'noise' is actually the model! Incidentally, this approach is the reverse of the regression techniques usually employed by astronomers. In most problems in astronomy, one tries to infer the parameters that describe the mean model and, if correlated noise is present, to marginalize over that noise. Here, the parameters describing the correlated noise are what we are interested in and the mean model is simply a straight line at $y = 0$.
[2] All computational times cited in this section are based on evaluations on a single core of a 2015 Macbook Pro, 3.1 GHz Intel Core i7.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

2.2.2 *Period*

We use an informed period prior, based on the ACF of the light curve. The ACF has proven to be very useful for measuring stellar rotation periods (McQuillan, Aigrain & Roberts 2012; McQuillan et al. 2013b; McQuillan, Mazeh & Aigrain 2014); however, the method has several shortcomings, most notably the inability to deliver uncertainties, but also the necessity of several heuristic choices, such as a time-scale on which to smooth the ACF, how to define a peak, whether the first or second peak gets selected, and what constitutes a secure detection. While this paper presents a rotation period inference method that avoids these shortcomings, it seems prudent to still use information available from the ACF. We thus use the ACF to define a prior on period, which can help the posterior sampling converge on the true period.

We do not attempt to decide which single peak in the ACF best represents the true rotation period, but rather we identify several *candidate* periods and define a weighting scheme in order to create a non-commital, though useful, multimodal prior. While this procedure does not avoid heuristic choices, we soften the potential impact of these choices because we are simply creating a *prior* for probabilistically justified inference rather attempting to identify a single correct period.

As another innovation beyond what ACF methods in the literature have presented, we also bandpass filter the light curves (using a 5th order Butterworth filter, as implemented in SCIPY) before calculating the autocorrelation. This suppresses power on time-scales shorter than a chosen minimum period $P_{\rm min}$ and longer than a chosen maximum $P_{\rm max}$, producing a cleaner autocorrelation signal than an unfiltered light curve.

We use the following procedure to construct a prior for rotation period given a light curve:

(i) For each value of P_i, where $i = \{1, 2, 4, 8, 16, 32, 64, 128\}$ d, we apply a bandpass filter to the light curve using $P_{\rm min} = 0.1$ d and $P_{\rm max} = P_i$. We then calculate the ACF of the filtered light curve out to a maximum lag of $2P_i$ and smooth it with a boxcar filter of width $P_i/10$.
(ii) We identify the time lag corresponding to the first peak of each of these ACFs, as well as the first peak's trough-to-peak height, creating a set of candidate periods T_i and heights h_i.
(iii) We assign a quality metric Q_i to each of these candidate periods, as follows. First, we model the ACF as a function of lag time t as a damped oscillator with fixed period T_i:

$$y = A\mathrm{e}^{-t/\tau}\cos\frac{2\pi t}{T_i}, \tag{7}$$

and find the best-fitting parameters A_i and τ_i by a non-linear least-squares minimization procedure – though enforcing a maximum value of $\tau = 20 \times P_i/T_i$ to avoid unreasonably large solutions of τ. We then define the following heuristic quality metric:

$$Q_i = \left(\frac{\tau_i}{T_i}\right)\left(\frac{N_i h_i}{R_i}\right), \tag{8}$$

where h_i is the height of the ACF peak at T_i, N_i is the length of the lag vector in the ACF (directly proportional to the maximum allowed period P_i), and R_i is the sum of squared residuals between the damped oscillator model and the actual ACF data. The idea behind this quality metric is to give a candidate period a higher score if

(a) it has many regular sinusoidal peaks, such that the decay time τ_i is long compared to the oscillation period T_i,
(b) the ACF peak height is high, and
(c) the damped oscillator model is a good fit (in a χ^2 sense) to the ACF, with extra bonus for being a good fit over more points (larger N_i, or longer P_i).

(iv) Given this set of candidate periods T_i and quality metrics Q_i, we finally construct a multimodal prior on the P parameter of the GP model as a weighted mixture of Gaussians:

$$p(\ln P) = \frac{\sum_i Q_i\,(0.9\mathcal{N}(\ln T_i, 0.2) + 0.05\mathcal{N}(\ln(T_i/2), 0.2) + 0.05\mathcal{N}(\ln(2T_i), 0.2))}{\sum_i Q_i}. \tag{9}$$

That is, in addition to taking the candidate periods themselves as mixture components, we also mix in twice and half each candidate period at a lower level, which compensates for the possibility that the first peak in the ACF may actually represent half or twice the actual rotation period. The period width of 0.2 in log space (corresponding to roughly 20 per cent uncertainty) is again a heuristic choice, balancing a healthy specificity with the desire to not have the results of the inference overly determined by the ACF prior.

Incidentally, while we use the procedure described here to create a prior on P which we use while inferring the parameters of the QP-GP model, this same procedure may also be used in the service of a rotation-period estimating procedure all on its own, perhaps being even more robust and accurate than the traditional ACF method. We leave exploration of this possibility to future work.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

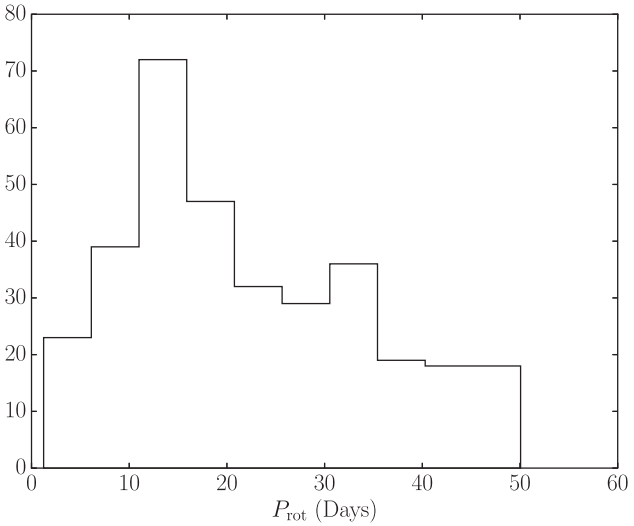

Figure 2. A histogram of the rotation periods used to generate the 333 simulated light curves in Aigrain et al. (2015).

2.3 Sampling

To sample the posterior in a way that is sensitive to potential multimodality, we use the EMCEE3[3] MCMC sampler. EMCEE3 is the successor to the EMCEE project (Foreman-Mackey et al. 2013) that includes a suite of ensemble MCMC proposals that can be combined to efficiently sample more distributions than the stretch move (Goodman & Weare 2010) in EMCEE. For this project, we use a weighted mixture of three proposals. First, we include a proposal based on the KOMBINE package[4] (Farr & Farr, in preparation), where a kernel density estimate (KDE) of the density represented by the complementary ensemble is used as the proposal for the other walkers. The other two proposals are a 'differential evolution (DE) MCMC' proposal (Ter Braak 2006; Nelson, Ford & Payne 2014) and the 'snooker' extension of DE (ter Braak & Vrugt 2008).

We initialise 500 walkers with random samples from the prior and use a weighted mixture of the KDE, DE and snooker proposals with weights of 0.4, 0.4 and 0.2, respectively. We run 50 steps of the sampler at a time, checking for convergence after each iteration, up to a maximum of 50 iterations. We declare convergence if the total effective chain length is at least 8× the maximum autocorrelation time. When convergence is achieved, we discard the first two autocorrelation lengths in the chain as a burn-in, and randomly choose 5000 samples as representative of the posterior. This fitting process takes several hours for a typical simulated light curve, though in some cases it can take 12 h or longer to converge.

3 PERFORMANCE AND COMPARISON TO LITERATURE: SIMULATED DATA

In order to test this new rotation period recovery method, we apply it to a set of simulated light curves and compare to the performance of established literature methods. Section 3.1 describes the simulated data we use; Section 3.2 demonstrates the performance of the QP-GP method; and Section 3.3 compares it to the performance of the LS periodogram and ACF methods.

3.1 Simulated light curves

We take our test data set from the Aigrain et al. (2015) 'hare and hounds' rotation period recovery experiment. These light curves result from placing dark, circular spots with slowly evolving size on the surface of bright, rotating spheres, ignoring limb-darkening effects. Aigrain et al. (2015) simulated one thousand such light curves to test the ability of participating teams to recover both stellar rotation periods and rotational shear (the amplitude of surface differential rotation). However, in this work, in order to focus on demonstrating reliable period recovery, we select only the 333 light curves without differential rotation, as differential rotation may produce additional scatter in the measured rotation periods.

Each of these light curve simulations uses a real *Kepler* long-cadence time array: one data point every 30 min over a 4 yr duration. 90 per cent of the rotation periods of the simulations come from a log-uniform distribution between 10 and 50 d, and 10 per cent from a log-uniform distribution between 1 and 10 d. Fig. 2 shows the distribution of solid-body rotation periods. The simulations also have a range of stellar inclination angles, activity levels, spot lifetimes and more (see Table 1). In order to preserve *Kepler* noise properties, Aigrain et al. (2015) add real *Kepler* light curves with no obvious astrophysical variability to the theoretical rotationally modulated light curves. Fig. 3 shows an example of a simulated light curve with a period of 20.8 d.

[3] https://github.com/dfm/emcee3
[4] https://github.com/bfarr/kombine

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

Table 1. Ranges and distributions of parameters used to simulate light curves in Aigrain et al. (2015).

Parameter	Range	Distribution
Rotation period ($P_{\rm rot}$)	10–50 d (90 per cent)	Log uniform
	1–10 d (10 per cent)	Log uniform
Activity cycle length	1–10 yr	Log uniform
Inclination	0°–90°	Uniform in $\sin^2 i$
Decay time-scale	(1–10) × $P_{\rm rot}$	Log uniform

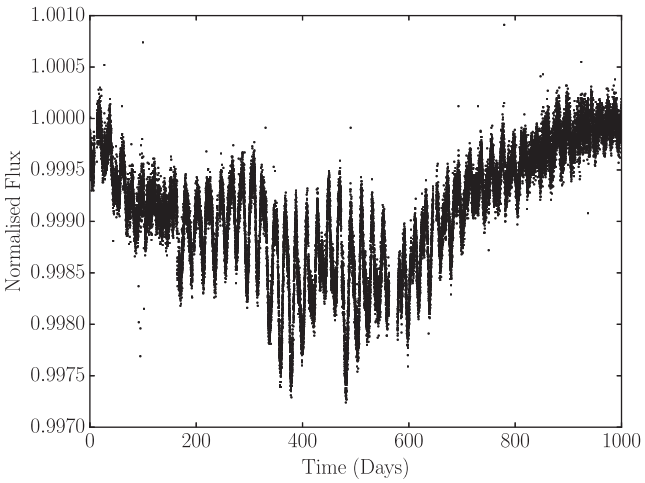

Figure 3. An example simulated light curve. This 'star' has a rotation period of 20.8 d.

Table 2. Priors and bounds on the natural logarithms of the GP model parameters.

Parameter	Prior	Bounds
$\ln A$	$\mathcal{N}(-13, 5.7)$	$(-20, 0)$
$\ln l$	$\mathcal{N}(7.2, 1.2)$	$(2, 20)$
$\ln \Gamma$	$\mathcal{N}(-2.3, 1.4)$	$(-10, 3)$
$\ln \sigma$	$\mathcal{N}(-17, 5)$	$(-20, 0)$
$\ln P$	Uniform/ACF based	$(\ln 0.5, \ln 100)$

3.2 Method performance

We apply the QP-GP inference method described in Section 2 to each of these 333 simulated light curves. As discussed in Section 2.2, reliable inference requires defining a useful set of priors on the non-period hyperparameters. For this simulation data set, we determined these by first running the method using very broad priors on all the non-period parameters (log flat between −20 and 20) and then inspecting the distribution of their posteriors for those cases that successfully recovered the true period. We also experimented with constraining the allowed ranges of the parameters after discovering that some regions of parameter space (such as large values of A and Γ and small values of l) tended to allow fits that ignored the desired periodicity. We list the final priors and bounds this process led us to adopt in Table 2. For the period prior, we tried two different methods: an uninformed (log flat) prior between 0.5 and 100 d, and an ACF-informed prior (Section 2.2.2).

Figs 4 and 5 summarize our results compared to the injected 'true' stellar rotation periods, for the uninformed and ACF-informed priors on period, respectively. The periods inferred using the GP method show good agreement with the true periods. To assess the performance of the QP-GP and other period recovery methods, we compute median absolute deviations (MADs) of the results, relative to the input periods. We also compute the median relative deviation (MRD) as a percentage and the root mean square (rms). These metrics are presented for the three different methods tested in this paper in Table 3. The informative and uninformative prior versions of the GP method have MRDs of 1.97 per cent and 2.62 per cent, respectively. The marginal posterior distributions of the QP kernel hyperparameters, e.g. simulated light curve in Fig. 3, are shown in Fig. 6.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

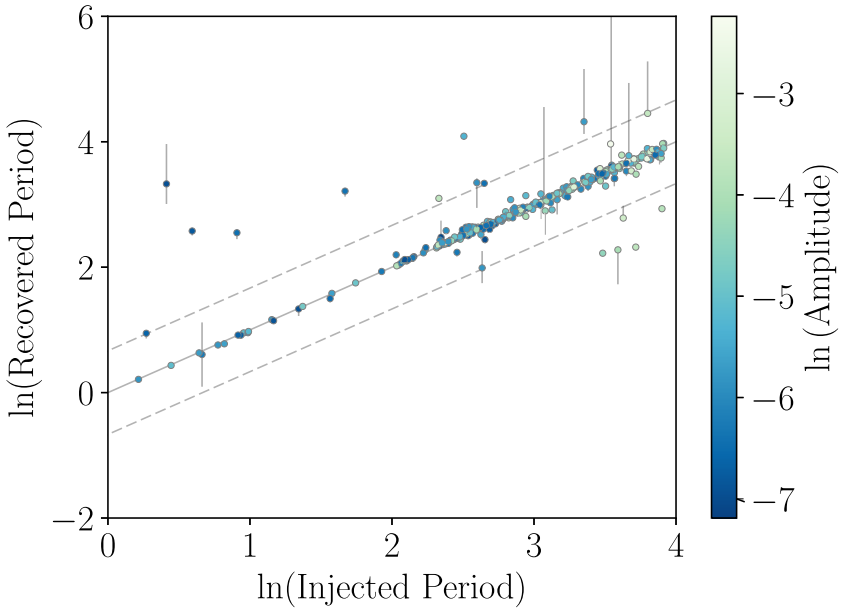

Figure 4. The 'true' rotation periods used to generate 333 simulated light curves versus the rotation periods measured using the GP technique with an informed, ACF-based prior. Points are coloured by the peak-to-peak amplitude of the light curve, as defined in Aigrain et al. (2015). Since the posterior PDFs of rotation periods are often non-Gaussian, the points plotted here are maximum a posteriori results. The uncertainties are the 16th and 84th percentiles. In many cases, the uncertainties are underestimated. The ACF-informed prior on rotation period used to generate these results is described in §2.2.

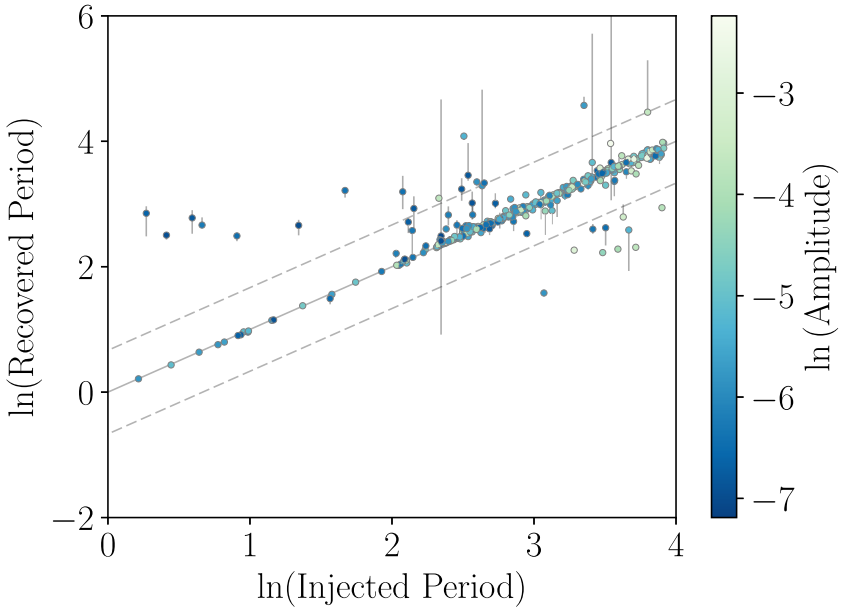

Figure 5. The 'true' rotation periods used to generate 333 simulated light curves versus the rotation periods measured using the GP technique with an uninformative prior. Points are coloured by the peak-to-peak amplitude of the light curve, as defined in Aigrain et al. (2015). Since the posterior PDFs of rotation periods are often non-Gaussian, the points plotted here are maximum a posteriori results. The uncertainties are the 16th and 84th percentiles. In many cases, the uncertainties are underestimated. An uninformative prior, flat in the natural log of the rotation period between 0.5 and 100 d was used to generate these results.

Table 3. Median absolute (MAD), median relative (MRD) deviations and rms for the LS periodogram, McQuillan et al. (2013a) ACF and GP period recovery methods.

Method	MAD (days)	MRD (per cent)	rms (d)
GP (ACF prior)	0.40	1.97	0.25
GP (uninformative prior)	0.48	2.62	0.44
ACF	0.37	1.75	4.58
LS periodogram	0.52	2.82	1.04

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

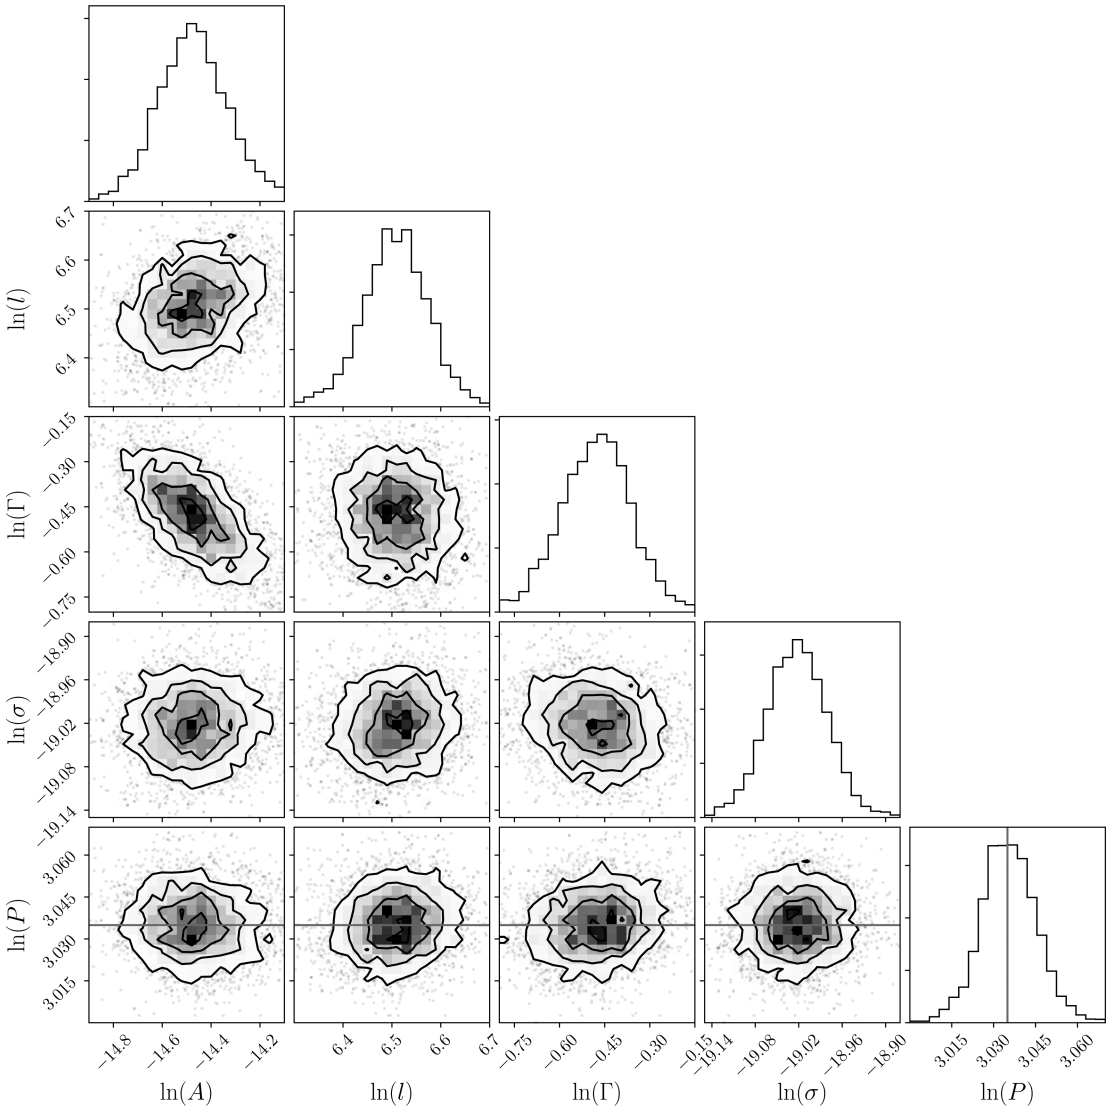

Figure 6. Marginal posterior PDFs of the QP-GP model parameters, fit to the simulated light curve in Fig. 3. The solid line in the period panel shows the injected period. This figure was made using `corner.py` (Foreman-Mackey 2016).

3.3 Comparison with literature methods

3.3.1 ACF

The ACF method has proven extremely useful for measuring rotation periods. The catalogue of rotation periods of *Kepler* stars provided in McQuillan, Aigrain & Mazeh (2013a) has been widely used by the community and has provided ground-breaking results for stellar and exoplanetary science. The method also performed well in the Aigrain et al. (2015) recovery experiment, producing a large number of accurate rotation period measurements (see e.g. their fig. 8). Another advantage is its fast implementation speed. However, because the ACF method is non-probabilistic, ACF-estimated rotation period uncertainties are poorly defined – a clear disadvantage. It also requires evenly spaced data; therefore, it is not applicable to ground-based light curves. Fig. 7 contains an example ACF of the light curve in Fig. 3.

We compare our results to those of the Tel Aviv team in Aigrain et al. (2015), who used the ACF method described in McQuillan et al. (2014) and Aigrain et al. (2015). This method involves the calculation of an ACF, a measure of the self-similarity of the light curve over a range of lags, which is then smoothed with a Gaussian kernel. Periodic variability in the light curve produces a peak in the ACF at the period of the signal and a decaying sequence of peaks at integer multiples of the period. McQuillan et al. (2014) first pre-processed the light curves by removing very long-term variations. They then calculated an ACF for each star and measured the positions of up to the first four peaks in the ACF. They fit a straight line to the periods of these peaks as a function of their integer multiple and adopted the gradient of that line as the rotation period. The uncertainty on the ACF period is the uncertainty on the best-fitting slope. Fig. 8 shows ACF-measured versus true rotation periods, with the $2n$ and $\frac{1}{2}n$ harmonic lines as dashed lines. The injected and recovered rotation periods generally agree well: the MRD of their periods, relative to the true rotation periods, is 1.75 per cent – slightly smaller than the results produced by the GP method.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

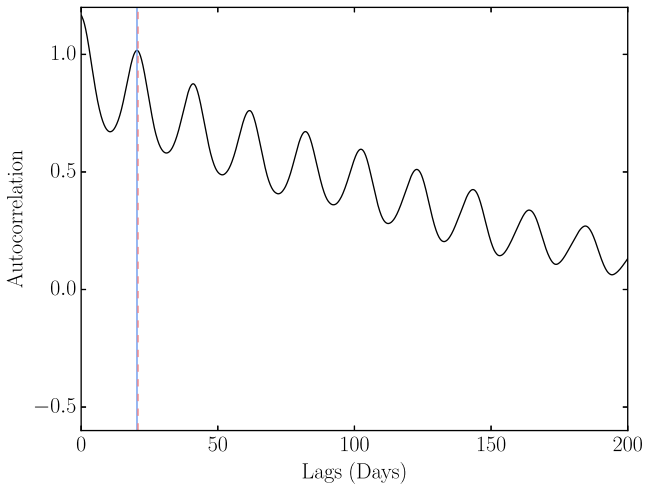

Figure 7. An ACF of the simulated light curve shown in Fig. 3. The vertical solid blue line shows the period measured using the ACF method (20.4 d) and the pink dashed line shows the period that was used to simulate the light curve (20.8 d).

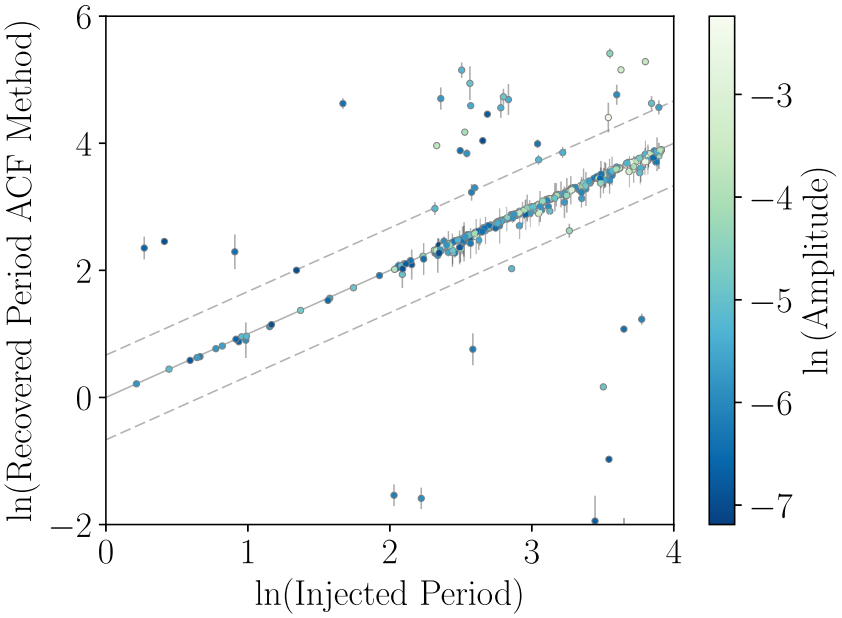

Figure 8. The 'true' rotation periods used to generate 333 simulated light curves versus the rotation periods measured using the ACF technique. Points are coloured by the peak-to-peak amplitude of the light curve, as defined in Aigrain et al. (2015). Several light curves have drastically over- and underestimated rotation periods.

However, their results contain a larger number of extreme outliers – they predict periods that are significantly different from the true period in more cases than the GP method. This is revealed by comparing the rms of the two sets of results since the rms is more sensitive to outliers. The rms of the Tel Aviv method results is 4.58 d and the rms of GP method results is 0.25 d.

3.3.2 LS periodogram

The Aigrain et al. (2015) Tel Aviv team were the only group to report periods for all 333 simulated stars, and are therefore the only team with which we compare results. The other participating groups chose to report only the period measurements in which they were confident and are likely to have omitted large outliers. In order, therefore, to test our results against another commonly used technique, we implemented a simple LS periodogram method.

We first applied a high-pass filter to the light curves, removing long-term trends that could produce large peaks at long periods in the periodogram. We used the third-order Butterworth filter in SCIPY with a cut-off of 35 d, attenuating signals with periods greater than this threshold. This filter was applied in frequency space, rather than period space, and full attenuation was attained at zero frequency. In other words, the level of attenuation increased smoothly with period, for periods above 35 d but no periods were fully suppressed. Experimentally, we found the 35 d cut-off removed the largest number of long period outliers while preserving shorter period signals. We also tried 50, 45,

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

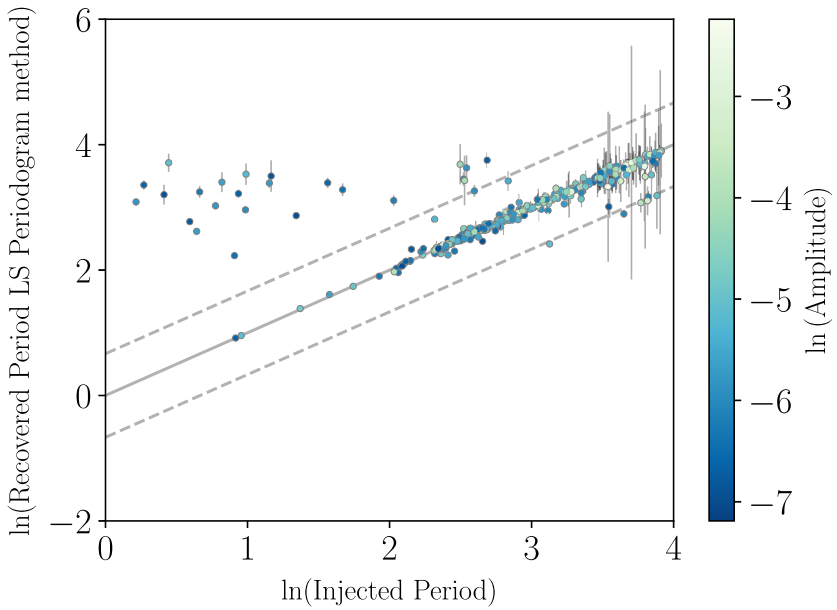

Figure 9. The 'true' rotation periods used to generate 333 simulated light curves versus the rotation periods measured using an LS periodogram technique. Points are coloured by the peak-to-peak amplitude of the light curve, as defined in Aigrain et al. (2015). In many cases, a large peak at a long period was present in the periodogram, producing a significant overestimate of the period.

40 and 30 d cut-offs, but the 35 d cut-off value minimized the rms and MAD of the LS periodogram results. For each simulated light curve, we computed an LS periodogram[5] over a grid of 10 000 periods, evenly spaced in frequency, between 1 and 100 d. We adopted the period of the highest peak in the periodogram as the measured rotation period. The uncertainties on the rotation periods were calculated using the following equation for the standard deviation of the frequency (Horne & Baliunas 1986; Kovacs 1981):

$$\sigma_\nu = \frac{3\pi\sigma_N}{2N^{1/2}TA}, \tag{10}$$

where A is the amplitude of the signal of highest power, σ_N is the variance of the time series, with the signal of highest power removed, N is the number of observations and T is the time span of the data. These formal uncertainties are only valid in the case that the noise is white, the data are evenly sampled and there is only one signal present. Since there are multiple signals present in these light curves, this formal uncertainty is a noisy estimate of the true uncertainty. Fig. 9 shows the resulting recovered rotation periods as a function of true period. The MRD of the periodogram-recovered periods is 2.82 per cent, slightly worse than the ACF and GP methods (see Table 3 for a side-by-side comparison with the other methods). The filter suppresses signals with periods greater than 35 d but doesn't eliminate them altogether, thus allowing high-amplitude frequencies to be detectable. For this reason, some erroneous long periods are recovered by the periodogram method. These results still show a marked improvement in comparison to periodograms calculated from unfiltered data. Simply adopting the highest peak of a periodogram calculated from the *unfiltered* simulated data resulted in an rms of 18.98 and relative MAD of 4.70.

4 REAL *Kepler* DATA

In order to test our rotation period inference method on real data, we apply it to a set of 275 *Kepler* Object of Interest (KOI) host stars, that had rotation periods previously measured by the ACF method by McQuillan et al. (2013a). We use the pipeline-corrected flux (`pdcsap_flux` column in the *Kepler* light-curve table), median-normalized and unit-subtracted, and mask out all known transiting planet candidate signals. As with the simulated light curves, we randomly subsample each light curve by a factor of 30 and split it into segments of about 300 points for the purposes of evaluating the likelihood. We also follow the same MCMC-fitting procedure as with the simulated data, using the ACF-based prior as before.

Initially, we also use the same priors on the hyperparameters for the KOIs as for the simulated light curves (Table 2). However, we found that $\ln l$ and $\ln \Gamma$ tended towards slightly different values than the simulations. We also found that the allowed hyperparameter range that we used for the simulations was too large for the KOI population, as maybe ∼15 per cent of the fits tended towards corners in the hyperparameter space, resulting in poor period measurements. As a result, after this initial test, we subsequently adjusted the priors and re-fit all the KOIs. The final priors and parameter ranges that we used are in Table 4. The posterior samples of the model parameters for the *Kepler* objects of interest are available online: https://zenodo.org/record/292340#.WKWpiBIrJE4.

[5] LS periodograms were calculated using the GATSPY PYTHON module: https://github.com/astroML/gatspy/tree/master/gatspy/periodic.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

Table 4. Priors and bounds on the natural logarithms of the GP model parameters, for *Kepler* light curves.

Parameter	Prior	Bounds
$\ln A$	$\mathcal{N}(-13, 5.7)$	$(-20, 0)$
$\ln l$	$\mathcal{N}(5.0, 1.2)$	$(2, 8)$
$\ln \Gamma$	$\mathcal{N}(1.9, 1.4)$	$(0, 3)$
$\ln \sigma$	$\mathcal{N}(-17, 5)$	$(-20, 0)$
$\ln P$	ACF based	$(\ln 0.5, \ln 100)$

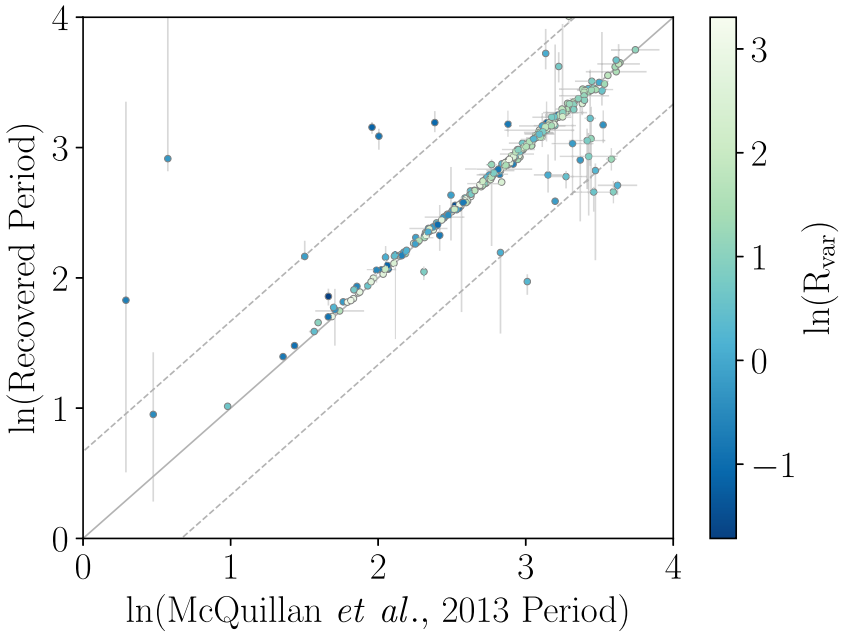

Figure 10. A comparison of our GP rotation period measurements to those of McQuillan et al. (2013a). The data points are coloured by the range of variability measured by McQuillan et al. (2013a), defined as the interval between the 5th and 95th percentiles of normalized flux per period bin in millimagnitudes.

Table 5. Rotation periods and *Kepler* input catalogue (KIC) properties for 1102 *Kepler* objects of interest. The full table is available online.

KOI number	Period (d)	KIC $T_{\rm eff}$ (K)	KIC log (g) (dex)	KIC [Fe/H] (dex)
1	25^{+23}_{-6}	5126^{+160}_{-138}	$4.61^{+0.03}_{-0.11}$	-0.12 ± 0.3
2	20 ± 4	5850 ± 50	4.46 ± 0.03	-0.15 ± 0.1
3	$28^{+3}_{-0.2}$	6350 ± 80	4.02 ± 0.01	0.26 ± 0.08
6	23^{+3}_{-15}	6225^{+114}_{-158}	$4.17^{+0.06}_{-0.05}$	$-0.04^{+0.1}_{-0.2}$
7	25 ± 2	5543 ± 79	4.08 ± 0.01	0.44 ± 0.1
...	...	...	...	...

Fig. 10 compares the periods inferred with the GP method to the ACF-based periods from McQuillan et al. (2013a) for the 275 overlapping KOIs. This comparison shows generally very good agreement, demonstrating that this method works not only on simulated data but also on real data – with the caveat that for any particular data set, some care is needed regarding the setting the priors and ranges for the GP hyperparameters. The new GP rotation periods calculated for 1102 KOIs are provided in an online table accompanying this paper. The first few lines of this table are shown in Table 5. Samples of the rotation period posterior PDFs are available online: https://doi.org/10.5281/zenodo.804440.

5 DISCUSSION

Upon examining the light curves that produce incorrect, long period measurements seen in the periodogram comparison figure (Fig. 9), it appears that the majority of these overpredictions are produced when the rotation period signal is highly non-sinusoidal. In these cases, the peak corresponding to the true period has a reduced amplitude and is not the highest peak in the periodogram. This can be understood intuitively as the LS periodogram fits a series of sine waves to the data; if the data are not sinusoidal, a sine wave will not produce a good

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

fit. There may be ways to further improve the precision of the periodogram method, for example modelling the detailed shapes of peaks in the periodogram instead of adopting the point closest to the derivative zero-point. However, we consider it likely that such alterations would produce only slight changes to the results, for example because peaks in the LS periodograms of *Kepler* light curves are usually narrow. It is difficult to assess the reason for the ACF incorrect recoveries because we do not have the individual ACFs computed by the Aigrain et al. (2015) Tel Aviv team; however, it seems that the ACF method is more sensitive than the LS periodogram method to noise. This particular implementation of the LS periodogram method is capable of detecting periodic signals at lower amplitudes that the ACF. We do not suggest that the LS periodogram method is better than the ACF method in general. It was made clear in Aigrain et al. (2015) that the performance of each of these rotation period detection algorithms depends sensitively on the implementation technique. The exact steps performed in the pre-processing of the light curves and the post-processing of the periodogram or ACF seem to affect the results strongly. This is also true for the GP method. Using a different MCMC sampler, different priors, different subsampling or a different kernel would all alter the results. Clearly, none of these rotation period measurement methods is perfect and they all require a lot of background engineering and tuning. This is simply due to inherent difficulty of inferring periods, or any parameter that produces a highly multimodal probability distribution, from noisy data.

The QP-GP inference method we test in this paper produces slightly more accurate rotation periods than the ACF or periodogram methods. Additionally, because it explores the posterior PDF with MCMC, the GP method produces probabilistically justified uncertainties. Unfortunately, these uncertainties still appear to be underestimated. Of the rotation periods recovered from the simulated light curves using the ACF-informed prior, only 25 per cent of the measured periods lie within 1σ of the true period, 50 per cent lie within 2σ and 66 per cent within 3σ. The largest outlier is 114σ away from the true value. The median uncertainty is underestimated by around a factor of 2. These numbers are similar for the uninformative prior results. We attribute these underestimated uncertainties to the model choice. Although more appropriate than a sinusoid, the QP-GP is still only an effective model. A perfect physical model of the star would produce more accurate uncertainties.

The QP kernel function represents a simple effective model of a stellar light curve. It can describe a wide range of QP behaviour and is relatively simple, with only a few hyperparameters. Nevertheless, it is still a somewhat arbitrary choice. Another valid choice would be a cosine function multiplied by a squared exponential, used by Brewer & Stello (2009) to model asteroseismic pulsations,

$$k_{i,j} = A \exp\left(-\frac{(x_i - x_j)^2}{2l^2}\right) \cos\left(\frac{\pi(x_i - x_j)}{P}\right). \tag{11}$$

This function produces a positive semidefinite matrix and has the P parameter of interest, but differs qualitatively from the QP function by allowing negative covariances. Is it realistic to allow negative covariances? In practice, the ACFs of *Kepler* light curves often go negative. However, many stars have two active regions on opposite hemispheres that produce two brightness dips per rotation, and it may be difficult to model such stars with a kernel that forces anticorrelation of points $\frac{1}{2}$ a period apart. It would be worthwhile to test this assumption and this alternative kernel function (and others) in the future.

Not all *Kepler* light curves show evidence of stellar rotation. In some cases, a star may have few or no active regions, be rotating pole-on, or be rotating so slowly that the *Kepler* data detrending pipeline removes any signal. In other cases, there may be another source of variability present in the light curve, generating a false period detection. These sources may be physical, e.g. binary star interactions, intra-pixel contamination from other astrophysical objects, pulsating variable stars, asteroseismic oscillations in giants and even stellar activity cycles. Identifying many of these astrophysical false positives falls outside the scope of this GP method (e.g. applying colour cuts to avoid giant contamination); however, for some, such as variable stars, they may have distinctive hyperparameters (e.g. long coherence time-scales) that identify them. Testing this may be an interesting follow-up study. As well as astrophysical contamination, instrumental sources may contribute to contaminating variability, e.g. temperature variations or pointing shifts of the *Kepler* spacecraft. These are unlikely to be periodic and, again, may produce unusual combinations of hyperparameters. We also hope to test this in the future.

In addition, we are continuing to develop several other aspects of this GP method:

(i) To use an updated method for calculating the GP likelihood and to drastically improve computational efficiency:`celerite`, a new method for rapid computation of the 'solve' and determinant calculations of covariance matrices necessary for GP inference, has recently been developed (Foreman-Mackey et al. 2017). `celerite` scales with the number of data points, $\mathcal{O}(N)$, rather than the $N \log(N)^2$ scaling of the `HODLR` algorithm used in this analysis. It requires a slightly different functional form for the kernel, as the increased speed is obtained by using mathematical properties of exponential functions. Kernels must be written as the sums of exponentials in order to be implemented in `celerite`; however, a close approximation to the kernel function used here is demonstrated to recover rotation periods in Foreman-Mackey et al. (2017).

(ii) To investigate more physically motivated priors: In this paper, we have only explored the physical interpretation of one parameter in the kernel function (P), which we related to stellar rotation periods, but others also warrant physically motivated priors. For example, the harmonic complexity parameter, Γ, defines the typical number of inflection points per function period; suitable priors may be informed by the typical properties of observed light curves. For the evolution time-scale, l, it is sensible to enforce $l > P$, such that functions have clear periodicity. Intuitively, if $l < P$, functions may evolve significantly over time-scales shorter than one period, so a function $f(t + P)$ would look significantly different to $f(t)$, in which case any claims of periodicity become dubious. More stringent criteria may be suggested by stellar astrophysical considerations, e.g. knowledge of typical active region evolution time-scales or persistence lifetimes (e.g. $l \sim 2P$). Finally, the amplitude hyperparameter A may be related to active region coverage. Unfortunately, because these hyperparameters control the covariance

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

structure of functions, their interpretation is not always straightforward; interpretation is further complicated by significant degeneracies between the hyperparameters. Therefore, we defer such considerations, along with strategies for efficient optimization or marginalization of hyperparameters, to a separate paper (currently in preparation).

(iii) To build in a noise model for *Kepler* data: Because it is a *generative* model of the data, the GP method models the rotation period at the same time as systematic noise features. One can then marginalize over the parameters of the noise model. This approach would be extremely advantageous for *Kepler* data, since long-term trends are often removed by the *Kepler* detrending pipeline. Marginalizing over the noise model at the same time as inferring the parameters of interest will insure that the periodic signal is preserved.

6 CONCLUSION

We have attempted to recover the rotation periods of 333 simulated *Kepler*-like light curves for solid-body rotators (Aigrain et al. 2015) using three different methods: a GP method, an ACF method and an LS periodogram method. We demonstrate that the GP method produces the most accurate rotation periods of the three techniques. Comparing our results with the results of the Aigrain et al. (2015) Tel Aviv team who implement the ACF method of McQuillan et al. (2013a), we find their results to be slightly more precise in a median sense than the GP method results, but with a larger number of significant outliers. Additionally, we measure the rotation periods of 275 *Kepler* objects of interest using the GP method and find that these results compare well to those measured previously by McQuillan et al. (2013a). The good agreement between these two sets of results demonstrates that the GP method works well on real *Kepler* data. Samples of the full set of 1102 KOI rotation period posterior PDFs can be found at https://doi.org/10.5281/zenodo.804440.

Unlike the ACF and LS periodogram methods, the GP method provides posterior PDF samples that can be used to estimate rotation periods uncertainties. In general, it provides more accurate rotation periods than either method and produces significantly fewer outliers. The GP method can also be applied to non-uniformly sampled light curves, unlike the ACF method.

Although an improvement on competing methods, the uncertainties are still underestimated by the GP method in many cases because it is an approximate model. A QP-GP is clearly a good model, as demonstrated here; however, it is still only an effective model, not an accurate physical one. It captures the posterior PDF of the periodic component of the covariance matrix of a time series, not the actual rotation period of a physical star. The marginalized posterior PDF does not therefore exactly reflect the probability of a rotation period, given the data. Similarly, the 16th and 84th percentile ranges, calculated from the posterior PDF of the QP-GP model, are not exactly the uncertainty of a physical rotation period. Although underestimated uncertainties are a drawback of this method, they are at least well motivated and well defined which is not the case for uncertainties calculated from the ACF and LS periodogram methods.

Although GP inference is slow as it requires a 'solve' and a matrix determinant calculation at every step of an MCMC, a new GP solver, `celerite` (Foreman-Mackey et al. 2017) shows promise for increasing the computational speed of GP rotation period inference by several orders of magnitude. We intend to integrate `celerite` into our code in the near future. How much computational time is too much is a question that depends on the available resources of each individual and their scientific priorities. GP period inference is trivially parallelizable, so if one has access to thousands of computer cores, one can infer rotation periods of thousands of stars in a few hours (perhaps just a few minutes with `celerite`). For single stars, GP rotation inference takes a reasonable amount of time and has already been used on exoplanet hosts (e.g. Haywood et al. 2014; Vanderburg et al. 2015).

The code used in this project is available at https://github.com/RuthAngus/GProtation/.

ACKNOWLEDGEMENTS

This work was supported by a grant from the Simons Foundation (#429961, R. A.). TDM is supported by NASA grant NNX14AE11G and acknowledges the hospitality of the Institute for Advanced Study, where part of this work was completed. VR thanks Merton College and the National Research Foundation of South Africa for financial support. We would like to thank the anonymous referee who's comments significantly improved this paper. Some of the data presented in this paper were obtained from the Mikulski Archive for Space Telescopes (MAST). STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. Support for MAST for non-*HST* data is provided by the NASA Office of Space Science via grant NNX09AF08G and by other grants and contracts. This paper includes data collected by the *Kepler* mission. Funding for the *Kepler* mission is provided by the NASA Science Mission directorate.

REFERENCES

Aigrain S. et al., 2015, MNRAS, 450, 3211
Aigrain S., Parviainen H., Pope B. J. S., 2016, MNRAS, 459, 2408
Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2015, IEEE Trans. Pattern Anal. Mach. Intell., 38
Angus R., Aigrain S., Foreman-Mackey D., McQuillan A., 2015, MNRAS, 450, 1787
Barclay T., Endl M., Huber D., Foreman-Mackey D., Cochran W. D., MacQueen P. J., Rowe J. F., Quintana E. V., 2015, ApJ, 800, 46
Barnes S. A., 2003, ApJ, 586, 464
Barnes S. A., 2007, ApJ, 669, 1167
Brewer B. J., Stello D., 2009, MNRAS, 395, 2226
Carter J. A., Winn J. N., 2010, Astrophysics Source Code Library, record ascl:1010.039
Czekala I., Andrews S. M., Mandel K. S., Hogg D. W., Green G. M., 2015, ApJ, 812, 128
Dawson R. I. et al., 2014, ApJ, 791, 89

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018

Dumusque X., Santos N. C., Udry S., Lovis C, Bonfils X, 2011, A&A, 527, A82
Evans T. M., Aigrain S., Gibson N., Barstow J. K., Amundsen D. S., Tremblin P., Mourier P., 2015, MNRAS, 451, 680
Foreman-Mackey D., 2016, J Open Source Softw., 24
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J, 2013, PASP, 125, 306
Foreman-Mackey D., Hoyer S., Bernhard J., Angus R., 2014a, 10.5281/zenodo.11989
Foreman-Mackey D., Hogg D. W., Morton T. D., 2014b, ApJ, 795, 64
Foreman-Mackey D., Agol E., Angus R., Ambikasaran S., 2017, AJ, 154, 220
García R. A. et al., 2014, A&A, 572, A34
Gibson N. P., Aigrain S., Roberts S., Evans T. M., Osborne M., Pont F., 2012, MNRAS, 419, 2683
Goodman J., Weare J., 2010, Commun. Appl. Math. Comput. Sci., 5, 1
Haywood R. D., 2015, PhD thesis, Univ. St Andrews
Haywood R. D. et al., 2014, MNRAS, 443, 2517
Hogg D. W., Myers A. D., Bovy J., 2010, ApJ, 725, 2166
Horne J. H., Baliunas S. L., 1986, ApJ, 302, 757
Jeffers S. V., Keller C. U., 2009, in Stempels E., ed., AIP Conf. Proc. Vol. 1094, 15th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun. AIP Publishing, New York, p. 664
Kawaler S. D., 1989, ApJ, 343, L65
Keeling C. D., Whorf T. P., 2004, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory
Kipping D. M., 2012, MNRAS, 427, 2487
Kovacs G., 1981, Ap&SS, 78, 175
Lanza A. F., Das Chagas M. L., De Medeiros J. R., 2014, A&A, 564, A50
Littlefair S. P., Burningham B., Helling C., 2017, MNRAS, 466, 4250
Lomb N. R., 1976, Ap&SS, 39, 447
McQuillan A., Aigrain S., Roberts S., 2012, in Griffin E., Hanisch R., Seaman R., eds, Proc. IAU Symp. 285, New Horizons in Time-Domain Astronomy. Cambridge Univ. Press, Cambridge, UK, p. 364
McQuillan A., Aigrain S., Mazeh T., 2013a, MNRAS, 432, 1203
McQuillan A., Mazeh T., Aigrain S., 2013b, ApJ, 775, L11
McQuillan A., Mazeh T., Aigrain S., 2014, ApJS, 211, 24
Nelson B., Ford E. B., Payne M. J., 2014, ApJS, 210, 11
Rajpaul V., Aigrain S., Osborne M. A., Reece S., Roberts S., 2015, MNRAS, 452, 2269
Rajpaul V., Aigrain S., Roberts S., 2016, MNRAS, 456, L6
Rasmussen C. E., Williams C. K. I., 2005, Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning). The MIT Press, Cambridge, USA
Reinhold T., Reiners A., Basri G., 2013, A&A, 560, A4
Rogers L. A., 2015, ApJ, 801, 41
Russell H. N., 1906, ApJ, 24, 1
Scargle J. D., 1982, ApJ, 263, 835
Skumanich A., 1972, ApJ, 171, 565
Ter Braak C. J., 2006, Stat. Comput., 16, 239
ter Braak C. J., Vrugt J. A., 2008, Stat. Comput., 18, 435
van Saders J. L., Ceillier T., Metcalfe T. S., Silva Aguirre V., Pinsonneault M. H., García R. A., Mathur S., Davies G. R., 2016, Nature, 529, 181
Vanderburg A. et al., 2015, ApJ, 800, 59
Wolfgang A., Rogers L. A., Ford E. B., 2016, ApJ, 825, 19

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

supp_data

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Downloaded from https://academic.oup.com/mnras/article-abstract/474/2/2094/4209242
by Said Business School user
on 24 January 2018