

K2SC: flexible systematics correction and detrending of *K2* light curves using Gaussian process regression

S. Aigrain,[★] H. Parviainen and B. J. S. Pope

Department of Physics, University of Oxford, Keble Road, Oxford OX3 9UU, UK

Accepted 2016 March 22. Received 2016 March 22; in original form 2016 February 8

ABSTRACT

We present $\kappa 2sc$ (*K2* Systematics Correction), a PYTHON pipeline to model instrumental systematics and astrophysical variability in light curves from the *K2* mission. $\kappa 2sc$ uses Gaussian Process regression to model position-dependent systematics and time-dependent variability simultaneously, enabling the user to remove both (e.g. for transit searches) or to remove systematics while preserving variability (for variability studies). For periodic variables, $\kappa 2sc$ automatically computes estimates of the period, amplitude and evolution time-scale of the variability. We apply $\kappa 2sc$ to publicly available *K2* data from Campaigns 3–5 showing that we obtain photometric precision approaching that of the original *Kepler* mission. We compare our results to other publicly available *K2* pipelines, showing that we obtain similar or better results, on average. We use transit injection and recovery tests to evaluate the impact of $\kappa 2sc$ on planetary transit searches in *K2* Pre-search Data Conditioning data, for planet-to-star radius ratios down to $R_p/R_* = 0.01$ and periods up to $P = 40$ d, and show that $\kappa 2sc$ significantly improves the ability to distinguish between true and false detections, particularly for small planets. $\kappa 2sc$ can be run automatically on many light curves, or manually tailored for specific objects such as pulsating stars or large amplitude eclipsing binaries. It can be run on ASCII and FITS light-curve files, regardless of their origin. Both the code and the processed light curves are publicly available, and we provide instructions for downloading and using them. The methodology used by $\kappa 2sc$ will be applicable to future transit search missions such as *TESS* and *PLATO*.

Key words: methods: data analysis – planetary systems – stars: rotation.

1 INTRODUCTION

NASA’s *Kepler* space mission brought about a revolution in high-precision, long baseline photometry, leading to major advances in exoplanet detection and characterization, asteroseismology, and the study of a wide range of stellar variability phenomena. After the failure of the second of its four reaction wheels in 2013, the original mission – by then in its extended phase – came to an end. The satellite was then re-purposed to survey four fields per year, located in or near the Ecliptic plane. This new mission, known as *K2*, delivers light curves of reduced photometric precision, owing to the reduced pointing accuracy, but offers unique opportunities to observe large numbers of bright ($V \leq 12$) Sun-like stars, low-mass stars, and nearby open clusters, amongst other key targets (Howell et al. 2014).

K2 data were initially released in the form of target pixel files (TPFs), which are time series of individual postage stamps centred

on each of the target stars. During the first year of operations, a number of teams developed pipelines to extract light curves from these TPFs (Armstrong et al. 2014, 2015; Vanderburg 2014; Vanderburg & Johnson 2014; Aigrain et al. 2015; Buyschaert et al. 2015; Huang et al. 2015; Libralato et al. 2016). Some of these are ‘all-purpose’ in the sense that they can be applied to any TPFs, some are specifically designed for particular kinds of stars (e.g. asteroseismic targets) or for the ‘superstamps’ made of many contiguous TPFs, which are used in *K2* to observe the crowded cores of dense open clusters. All of them include a light-curve extraction step, which uses either aperture photometry (with circular or pixelated masks, whose positions can be static throughout an observation campaign or adjusted to follow the star’s position as the pointing varies) or point spread function fitting. All these pipelines also include a step to correct the systematic flux variations induced by the variations in the satellite’s pointing (sometimes known as detrending). Almost without exception, the systematics correction in the aforementioned pipelines is done on an object-by-object basis, by modelling the dependence of the star’s flux on its position. This represents a major departure from the approach most commonly used during the *Kepler* mission,

[★] E-mail: Suzanne.Aigrain@astro.ox.ac.uk

as implemented in the PDC^1 and later PDC-MAP^2 components of the standard *Kepler* pipeline, where ‘co-trending basis vectors’ (CBVs) were constructed by combining many observed light curves, and each light curve was then decomposed into a linear combination of these CBVs plus a residual term including the star’s intrinsic variations and random noise (Smith et al. 2012; Stumpe et al. 2012, 2014; Van Cleve et al. 2015).

From Campaign 3 onwards, the *K2* mission also started releasing light-curve files (LCFs) extracted from the TPFs using a slightly modified version of the standard *Kepler* Science Operations Center (KSOC) pipeline. These LCFs include both the simple aperture photometry (SAP) fluxes and the PDC fluxes, but the latter typically contain significant residual systematics, indicating that the PDC correction alone is not sufficient for *K2*. This motivated us to develop a somewhat improved and stand-alone version of the star-by-star systematics-correction step initially developed in Aigrain et al. (2015, hereafter Paper I), which can be applied to the SAP or PDC light curves – but also, in principle, to the light curves produced by any of the aforementioned photometry pipelines. This paper presents this systematics correction tool, which we call $\kappa 2\text{SC}$ (*K2* systematics correction), including details of its implementation, photometric performance evaluation and signal injection and recovery tests, and links to the code (which is open source and freely available) as well as to all the SAP and PDC data we have processed so far (all the individual long-cadence TPFs from Campaigns 3 to 5).

Our approach consists in modelling the observed flux as a Gaussian Process (GP) with three additive components: the first depends on the star’s position, and represents the pointing-induced systematics; the second depends on time, and represents the star’s intrinsic variability as well as any other systematics not dependent on position; and the third represents white noise. The first component can then be subtracted to leave a light curve corrected for position-dependent systematics. The use of GP regression gives our model flexibility and robustness, and avoids having to specify an arbitrary functional form for the variations we seek to model. It is implemented in a Bayesian framework, which enables us to incorporate relevant physical information, for example, in the form of priors on the length-scales of the various components, and to propagate the uncertainties associated with the systematics correction through to subsequent analysis. Our decision to model the star’s intrinsic variations explicitly alongside the systematics was initially motivated by the desire to be able to correct systematics in variable stars without affecting the variability itself, but as we will show it also improves the systematics correction even in the case of relatively quiet stars. As a bonus, the time component can also be subtracted to produce a light curve corrected for both systematics and stellar variability, which can be used, for example, to search for planetary transits.

Several improvements have been made to our systematics correction method since Paper I. First, while we originally modelled the systematics as a function of the satellite roll-angle only, we now model them as a function of each star’s 2D position (x and y). This gives better results for full *K2* Campaigns where the position variations include a significant drift as well as quasi-periodic roll-angle variations on ~ 6 -h time-scales. Furthermore, the same roll-angle variations can correspond to very different actual changes in position at different locations in the field of view (FOV), so using the star’s own position gives a better representation of the inter-

intra-pixel sensitivity variations which give rise to the systematics. We also introduced priors on the parameters of the GP model (length-scales and amplitudes for the various components) rather than performing a simple likelihood optimization within bounds. Finally, in addition to the standard (squared-exponential) kernel used to represent the stellar variability, which is suitable for smooth, aperiodic variations, we also implemented a quasi-periodic kernel which is better able to reproduce the behaviour of spotted stars and some classes of pulsators, along with a simple prescription for determining when this alternative kernel should be used.

Our method is described in detail in Section 2. We use two approaches to evaluate the performance of our pipeline. The first is to measure the light-curve scatter on transit time-scales, using a proxy measure of the Combined Differential Photometric Precision (CDPP); these tests are reported in Section 3. The second is to inject simulated planetary transit signals into the raw light curves and test our ability to recover them. Those tests are reported in Section 4. The $\kappa 2\text{SC}$ package and data from Campaigns 3 to 5 are publicly available, and Section 5.1 describes how to access and use them. We conclude and outline plans for future development in Section 6.

2 METHOD

2.1 The basics

We use GP regression to model the instrumental systematics, astrophysical variability and white noise in each light curve. A full description of GP regression is beyond the scope of this paper; we refer the interested reader to Rassmussen & Williams (2006) for a textbook introduction, and Gibson et al. (2012) and Aigrain, Pont & Zucker (2012), for examples of GP regression applied to exoplanet data sets. Formally, a GP is a stochastic process, such that the joint probability distribution over any collection of observations of this process is a multivariate Gaussian. The covariance matrix of this distribution is specified through a covariance function, which defines the covariance between pairs of observations as a function of a collection of input variables. This sets up a probability distribution over functions, which all share the desired covariance properties. Using well-known conditioning and marginalization identities for Gaussian distributions, it is then straight forward (and analytic) to condition this prior on the available data, and to evaluate the likelihood of the model, without having to specify a mathematical expression for the function directly. Standard optimization and model comparison methods can be used to fit for the parameters of the covariance function and/or to compare different covariance functions. Importantly, if the covariance function contains several, additive terms, it is possible to separate the contribution of the different terms. We exploit this property here to separate the position-dependent systematics from the time-dependent astrophysical signal.

In the original version of our pipeline, described in Paper I, we used a single coordinate, representing the global roll-angle variations of the satellite, as the input driving the systematics component. Instead, the present version uses the 2D position of each star. When processing the KSOC light curves, our input variables representing position are the `POS_CORR1` and `POS_CORR2` columns of the Mikulski Archive for Space Telescopes (MAST) light curves, which represent the deviation from the star’s nominal position on the detector, as estimated by the KSOC pipeline from the satellite motion polynomials.

Each light curve is an array of flux measurements $f = \{f_1, f_2, \dots, f_n\}$. For simplicity, we will assume these have been

¹ ‘PDC’ stands for ‘pre-search data conditioning’.

² ‘MAP’ stands for ‘maximum a posteriori’.

normalized by dividing by the median and subtracting unity. The joint distribution over the flux measurements is assumed to be a multivariate Gaussian with zero mean and covariance matrix \mathbf{K} :

$$\mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}). \quad (1)$$

The elements of the covariance matrix are given by

$$\begin{aligned} K_{ij} &= k(\mathbf{X}_i, \mathbf{X}_j) \\ &= k_{xy}(x_i, y_i, x_j, y_j) + k_t(t_i, t_j) + k_w(i, j), \end{aligned} \quad (2)$$

where $\mathbf{X}_i \equiv \{t_i, x_i, y_i\}$ and the three terms on the right-hand side represent, respectively, position-dependent systematics, the intrinsic variability of the star, and white noise. The position component uses a squared exponential covariance function with amplitude A_{xy} and separate inverse length-scales η_x and η_y :

$$k_{xy}(x_i, y_i, x_j, y_j) = A_{xy} \exp[-\eta_x(x_i - x_j)^2 - \eta_y(y_i - y_j)^2].$$

This gives rise to smooth variations with characteristic amplitude A_{xy} and length-scales $1/\eta_x$ and $1/\eta_y$ in the x - and y -directions, respectively. (We also tried using a single length-scale, i.e. a radial covariance function, but obtained slightly better results with separate length-scales for x and y .) By default, the time component also uses a squared exponential covariance function with amplitude A_t and inverse length-scale η_t :

$$k_t(t_i, t_j) = A_t \exp[-\eta_t(t_i - t_j)^2].$$

This likewise gives smooth variations with characteristic amplitude A_t and time-scale $1/\eta_t$. This works well so long as the star's intrinsic variability does not occur on time-scales similar to the pointing variations. In Section 2.4, we introduce a prescription for automatically switching to a more adequate covariance function for periodic and quasi-periodic variables. Finally, the white noise component is simply

$$k_w(i, j) = \sigma^2 \delta_{ij},$$

where δ_{ij} is the Kronecker delta function.

The log-likelihood of the model is now

$$\log p(\mathbf{f}|\theta) = -\frac{1}{2}N \log(2\pi) - \frac{1}{2} \log(|\mathbf{K}|) - \mathbf{f}^T \mathbf{K}^{-1} \mathbf{f},$$

where $\theta \equiv A_{xy}, \eta_x, \eta_y, A_t, \eta_t, \sigma$ are the parameters of the covariance function, $|\mathbf{K}|$ is the determinant of the covariance matrix, and \mathbf{f}^T is the transpose of \mathbf{f} . The GP kernels and log-likelihood are implemented using the `GEORGE` package³ (Ambikasaran et al. 2014).

We define priors over each of the parameters of the covariance function:

- (i) A_{xy} and A_t : uniform prior between -7 and 1 in \log_{10} normalized flux (that is, flux divided by its median);
- (ii) σ : uniform prior between -6 and 0 in normalized \log_{10} flux;
- (iii) η_x and η_y : truncated normal prior with mean 17 pixel^{-1} , standard deviation 8 pixel^{-1} , minimum 0 pixel^{-1} , and maximum 70 pixel^{-1} ;
- (iv) η_t : truncated log-normal prior with mean 0.25 d^{-1} , standard deviation 1.25 dex , and upper boundary at 2 d^{-1} to prevent overfitting the noise.

We learn the GP hyperparameters from the data by finding the posterior density maximum using a global optimization step followed by local optimization. First, we use a differential evolution (DE)

global optimizer (Storn & Price 1997) adapted from `PyDE`⁴ to explore the full range of parameter space and generate a population of parameter vectors clumped close to the global posterior maximum. The DE starting population is initialized to uniformly sample a sub-volume of the parameter space allowed by the priors (this is because the uniform prior high boundaries are much higher than what would be expected for any realistic case), but the population is allowed to expand to explore outside the initial boundaries. After the DE step, the parameter vector with the highest posterior value is chosen as a starting point for local optimization using Powell's local optimization method (Powell 1964, as implemented in `SCIPY`).

For a given set of covariance parameters, the predictive distribution for the model conditioned on the data, evaluated at any set of inputs $\mathbf{X}_* = \{t_*, x_*, y_*\}$, is a Gaussian with mean

$$\bar{\mathbf{f}}_* = \mathbf{K}_* \mathbf{K}^{-1} \mathbf{f}$$

and variance

$$\text{var}(f_*) = \mathbf{K}_{**} - \mathbf{K}_* \mathbf{K}^{-1} \mathbf{K}_*^T,$$

where $\mathbf{K}_* \equiv \{k(\mathbf{X}_*, \mathbf{X}_1), k(\mathbf{X}_*, \mathbf{X}_2), \dots, k(\mathbf{X}_*, \mathbf{X}_n)\}$ and $\mathbf{K}_{**} \equiv k(\mathbf{X}_*, \mathbf{X}_*)$. Importantly, we can evaluate the predictive mean for the different components of the model separately

$$\bar{f}_{*,t} = \mathbf{K}_{*,t} \mathbf{K}^{-1} \mathbf{f},$$

where $\mathbf{K}_{*,t} \equiv \{k_t(t_*, t_1), k_t(t_*, t_2), \dots, k_t(t_*, t_n)\}$, and similarly

$$\bar{f}_{*,xy} = \mathbf{K}_{*,xy} \mathbf{K}^{-1} \mathbf{f},$$

where $\mathbf{K}_{*,xy} \equiv \{k_{xy}(x_*, y_*, x_1, y_1), k_{xy}(x_*, y_*, x_2, y_2), \dots, k_{xy}(x_*, y_*, x_n, y_n)\}$.

We can thus remove the position-dependent systematics from the raw light curve, while preserving intrinsic variability, by evaluating $\bar{f}_{*,xy}$ at each cadence, and subtracting it from the original. Furthermore, it is also straightforward to evaluate the variance of the predictive mean for the position component alone, and hence to obtain a robust uncertainty estimate for each point in the corrected light curve.

We also compute and store the full model prediction ($\bar{\mathbf{f}}_*$ evaluated at each cadence), as this can be subtracted from the raw data to obtain a version of the light curve where both systematics and the smooth component of any intrinsic variability have been removed: this can in principle be useful in searching for short-duration events such as eclipses, planetary transits and flares.

Figs 1 and 2 show examples of the light curve before and after processing, and the time- and position-dependence of the model, for a typical star observed in *K2* Campaign 5. In this case, the 'raw' light curve was the `SAP` flux. When working with LCFs produced by the standard *Kepler* pipeline, we use the position of the star predicted from the motion polynomials of the satellite (stored in the `POS_CORR1` and `POS_CORR2` columns), which should be more precise and less sensitive to contamination effects than the centroid-based estimates (stored in the `MOM_CENTR1` and `MOM_CENTR2` columns).

2.2 Outlier detection

If left untreated, outliers can be a significant problem: the model seeks to explain them by adopting an unphysically short length-scale for the time component. To address this, we use an initial

³ See <http://dan.iel.fm/george>

⁴ <https://github.com/hpparvi/PyDE>, DOI:10.5281/zenodo.45602




Figure 1. Example: Campaign 5 *K2* light curve before and after processing. From top to bottom: x and y position (black and red points, respectively), raw flux, flux corrected for systematics, flux corrected for intrinsic variations (showing systematics only), and residuals of the full model (after subtracting both the time and the position components). In the bottom four panels, black points show the data, the red curve shows the model, the pink shaded areas show the 95 per cent confidence interval about the model, and the points circled in red are identified as outliers and excluded from the model calculation. The vertical dotted line marks the point at which the direction of the roll-angle variations reverse, the position-dependent variations are treated separately either side of this line (see Section 2.3).

outlier detection step. We carry out an initial detrending using a default set of hyperparameters based on the complete hyperparameter population of Campaign 4 (see Section 2.5 for a discussion of the hyperparameter distributions), and flag any points deviating more than 5σ above or below of the detrended light curve as outliers. The kernel hyperparameter optimization is then carried out using a subset of the remaining points, and the predictive distribution is computed again using the best-fitting parameters. Points located more than 5σ above or below this prediction are again flagged, and a new, final prediction is generated. This final prediction uses only points which were not flagged as outliers, but it is computed for all data points, including outliers. For example, deep transits and stellar eclipses are often flagged as outliers, but it is still desirable to compute a corrected light curve including these events.

Despite the relatively sophisticated outlier treatment described above, eclipses are not always flagged as outliers. Specifically, if there is significant out-of-eclipse variability on time-scales similar to the eclipse duration, the length-scale of the time component converges to a value that is short compared to the eclipse duration. Consequently, the time component of the model attempts to fit the eclipses as well as the out-of-eclipse variability, but generally struggles to reproduce the sharp ingress and egress of the eclipses. To illustrate this, Fig. 3 shows light-curve segments for two eclipsing binaries from *K2* Campaign 5. The first has short-duration

eclipses and little out-of-eclipse variability, and the eclipses are duly flagged as outliers. The second has significant variability and longer eclipses, and the in-eclipse points are flagged as outliers. In such a case, not only would removing the time component of the model remove most of the eclipse signal, but the systematics correction is also less satisfactory during the eclipses. To address this issue, we included an option in $\kappa 2sc$ to manually force in-eclipse points to be treated as outliers for individual objects, by specifying the ephemeris and eclipse duration.

2.3 Break-points

Visual examination of *K2* light curves before and after modelling them using the method described above soon revealed that the behaviour of the systematics changes qualitatively at one or two points during each campaign. This is clearly visible in the light curve shown in Fig. 4, for example (which also illustrates how we deal with periodic variable stars, as discussed in Section 2.4). These points correspond to reversals of the direction of the roll-angle variations: the net torque due to solar radiation pressure pushes the satellite first one way, then the other as the campaign progresses and the orientation of *Kepler* changes relative to the Sun. It is not entirely clear why this leads to different behaviours in the systematics, it may be due to small changes in the star's position during each exposure




Figure 2. Position-dependence of the flux after removing the intrinsic variations (left) and the position-dependent component of the model (right), for the light curve shown in Fig. 1. The top and bottom panels show data taken before and after the break-point (dotted vertical line on Fig. 1). The data and model agree closely, and the differences between the behaviour before and after the break-point are noticeable. Note also that the variations are clearly 2D, so a 1D model for the position dependence would not work as well. Finally, in the bottom row we see a hint of periodicity in the x -axis, with a period of 1 pixel, which is what one would expect if the main source of systematics is to intra-pixel sensitivity variations.




Figure 3. Segments of the systematics-corrected pDC light curves for two eclipsing binaries from K2 Campaign 5, illustrating the successful (top) and unsuccessful (bottom) identification of in-eclipse points as outliers. The symbols and lines are the same as for Fig. 1.

(30 min for long-cadence data). We modified the position-dependent term in the covariance function to include a break-point each time the roll-angle variations change direction. This is implemented by multiplying the original covariance matrix with a binary mask which takes the value of unity if both points belong to the same segment (between two reversals of the roll-angle drift) and zero otherwise. The break-points between segments are the same for all objects in a given K2 Campaign, and are determined by visual inspection of the

position variations for a few dozen light curves spanning the FOV. We note that changing the precise timing of the break points (up to a day or two) does not affect the results significantly.

2.4 Handling variable stars

The model described in Section 2.1 performs well for stars with moderate intrinsic variability occurring on time-scales considerably longer than the characteristic time-scale of the K2 roll angle variations (approx. 6 h). In those cases, the two components of the GP adequately separate the intrinsic variability from the position-dependent systematics. However, the correction is less successful for strong variables such as classical pulsators or active, rapidly rotating stars. In those cases, as already noted in Paper I, A_t and A_{xy} tend to shrink to zero, and both systematics and intrinsic variability are typically absorbed into the white noise, which becomes much larger than normal (for a given star brightness). The correction is thus ‘conservative’, meaning that it does not remove true variability, but it also leaves much of the systematics unaffected.

In Paper I, we noted that this problem could be overcome on a case-by-case basis by altering the initial guess for η_t , or by implementing a more explicit model of the variability. This would be suitable for relatively rare kinds of variables, such as pulsating stars. However, one of the strengths of K2 is its ability to observe young open clusters, whose members are typically active and rapidly rotating. These are too numerous to be treated manually. We have therefore implemented an automated procedure for identifying and handling variable stars that display a clear periodicity. First, we compute the Lomb–Scargle periodogram of the raw light curve (Lomb 1976; Press et al. 2007) in the period range 0.05–20 d. If the false alarm probability of the periodogram maximum is lower than a given threshold value (by default 10^{-50}), we replace the time component of the GP model with the following, quasi-periodic covariance function:

$$k_{\text{rot}}(t_i, t_j) = A_t \exp \left[-\Gamma \sin^2 \left(\frac{\pi |t_i - t_j|}{P} \right) - \frac{(t_i - t_j)^2}{L_c^2} \right],$$

where P is the period, Γ the inverse length-scale of the periodic component of the variations, and L_c the evolutionary time-scale of the variations. P is initially set to the period of the periodogram peak, and the evolutionary time-scale to 10 times that value. The covariance parameters are then refined using the same procedure as for the non-periodic case. This gives significantly improved results for spotted stars and pulsating stars with pulsation periods of a day or more. As a by-product, this procedure yields estimates of the dominant period of the stellar variability and of its characteristic evolutionary time-scale, which are stored in the headers of the corrected light-curve files. Fig. 4 shows an example light curve for a spotted star before and after correction.

2.5 Hyperparameter distributions

Fig. 5 shows the distributions of the final (best-fitting) hyperparameters for Campaigns 3 to 6, for the non-periodic and quasi-periodic cases, respectively. The distributions are broadly consistent between the different campaigns, which implies that the noise, systematics and variability properties of the light curves do not change significantly from one campaign to the next. This provides an a-posteriori justification of our choice to use the medians of the distributions from Campaign 4 (which was the first campaign we processed in full) as the default values for the initial detrending performed prior to identifying outliers (see Section 2.2).




Figure 4. Same as Fig. 1, but for a light curve displaying quasi-periodic variations.

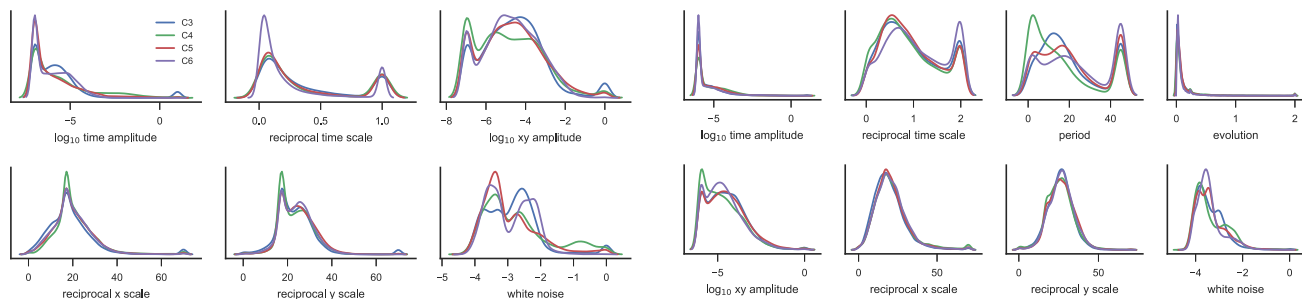


Figure 5. Final hyperparameter distributions for Campaigns 3 to 6, for the non-periodic cases (left) and the quasi-periodic cases (right).

The distributions for the white noise term show more variation between campaigns, but this is due to the different magnitude distributions of the targets. In the quasi-periodic cases, the final periods do not necessarily match the initial guess taken from the Lomb–Scargle results. The secondary peak in the period distribution at around 30 d, corresponds to cases where the period and evolutionary time-scales are similar. In such cases, the model reverts to a random, rather than clearly periodic behaviour. Focusing on periods below 30 d, we also note that the distribution for Campaign 4 peaks at shorter periods, which may be a result of the larger fraction of young (Pleiades and Hyades) stars observed in this campaign.

We also examined the hyperparameter distributions as a function of location in the FOV, and noted that the hyperparameters associated with the systematics component of the model depend

somewhat on distance from the satellite boresight. Incorporating this information in the initial guesses may improve the detrending, and will be considered in the future.

3 PHOTOMETRIC PERFORMANCE FOR K2 DATA

The CDPP has become a de-facto standard for evaluating the photometric precision of *Kepler* and *K2* light curves. It is formally defined as the inverse signal-to-noise ratio (SNR) of a reference transit signal of the corresponding duration, in parts per million (ppm), and can be interpreted as the depth of a transit of the given duration to exhibit an SNR of 1 (Christiansen et al. 2012). The KSOC transit search pipeline systematically evaluates the CDPP on 3, 6.5 and 12-h time-scales. Since we do not have access to this pipeline, we




Figure 6. 6.5-h CDPP versus *Kepler* magnitude for K2 Campaign 5, starting from the *SAP* or the *PDC* light curves (top and bottom, respectively). From left to right: *Kepler* pipeline CDPP estimates (these are available for *PDC* data only), and our CDPP estimates for the original data, the data corrected for systematics, and the data corrected for systematics and stellar variability. Objects identified in the reduced proper motion diagram (Fig. 7) as likely giants are shown in red.

can only evaluate a proxy measurement of the CDPP. To do this, we follow the same procedure as Gilliland et al. (2011): trends on time-scales longer than 2 d are removed using a Savitzky–Golay filter, then 5σ clipping is used to remove large outliers. We then compute the average flux in consecutive 13-point bins, keeping only the bins with >7 valid data points. The 6-h CDPP proxy is then estimated as the standard deviation of the bin-averaged fluxes. For each light curve processed with $\kappa 2sc$, this evaluation is performed on the input light curve, the light curve after removing systematics (i.e. subtracting the position-dependent component of the GP model), and the light curve after removing both systematics and stellar variability (i.e. subtracting also the time-dependent component of the GP model).

Fig. 6 shows the 6-h CDPP proxy for all light curves available from MAST for K2 Campaign 5 (data release 7, does not include superstamps or short-cadence targets). The data are shown as a function of *Kepler* magnitude (Kp), as provided in the EPIC catalogue (Huber et al. 2016). Only objects for which the Kp value included in the EPIC catalogue was derived from the Sloan g , r and i magnitudes are shown (as opposed to values derived from other magnitude systems such as B and V or 2MASS JHK). In the first column, we show for reference the 6-h CDPP values estimated by the KSOC pipeline (these are available for the *PDC* data only). The remaining three columns show our own CDPP proxy estimates, for the ‘raw’, systematics-corrected and systematics- and variability-corrected data (second, third and fourth columns, respectively). Note that we have also processed the data for Campaigns 3 and 4, and the results are similar.

The first and second panels in the bottom rows should look similar if our CDPP proxy is accurate; it does for the majority of the points, with the notable exception of those shown in red. These




Figure 7. Reduced proper motion diagram for K2 Campaign 5 targets. The black points follow the locus of Galactic disc main-sequence stars. The red points, corresponding to stars with red colours and low proper motion for their magnitude, are likely red giants.

are objects identified as likely giants in a reduced proper motion diagram (see Fig. 7), and are discussed further in Section 3.1. Our CDPP proxy estimates for the *SAP* data are considerably higher and more scattered, as expected since the *PDC* step partially removes both




Figure 8. 6.5-h CDPP versus *Kepler* magnitude for *K2* Campaign 5, excluding the giants and focusing on the brighter part of the magnitude range. The horizontal magenta lines show the median CDPP in 1-mag bins. For comparison, the cyan line shows the median CDPP obtained for stars with $11.5 \leq K_p \leq 12.5$ observed during the original *Kepler* mission (Gilliland et al. 2011).

the *K2* systematics and stellar variability on time-scales of days to weeks (Smith et al. 2012; Stumpe et al. 2012, 2014).

Giants aside, the CDPP proxy values for systematics-corrected data (third column) are considerably lower on average, and follow a tighter relation with K_p magnitude, than those for the raw light curves. This shows that our κ_{2SC} pipeline yields a significant improvement over the basic data products, and holds whether the *SAP* or the *PDC* data are used as the starting point, though the *PDC* light curves yield slightly lower CDPP values. Removing the time component of our model (final column) reduces the number of outliers lying above the main relation (variable stars) but otherwise does not change the diagram significantly.

In Fig. 8, we take a closer look at the photometric performance for dwarf stars with $K_p \leq 14.5$. The lower envelope of the distribution ranges from <20 ppm (for $K_p \sim 10$) to about 70 ppm (for $K_p \sim 14$), which is comparable with the performance recorded during the original *Kepler* mission. This implies that, for a significant fraction of the stars, we are able to remove the additional systematics caused by the roll-angle variations of the satellite to a level where they are below other sources of noise. We have also computed the median CDPP values in 1-mag bins (magenta lines). For the bin $11.5 \leq K_p \leq 12.5$, the median value is 38 ppm. This is only 50 per cent higher than the value of 25 ppm median CDPP derived in exactly the same manner, for the same magnitude range, for light curves from the original *Kepler* mission (Gilliland et al. 2011).

3.1 Giant stars

Red giants display excess variability on few hours time-scales due to stochastically excited pulsations. The *KSOC* transit-search pipeline includes a harmonic-filtering step which is designed to filter out these oscillations (Jenkins et al. 2010; Tenenbaum et al. 2012), and hence yields lower CDPP estimates for giants than our proxy values, which are based on the *PDC* data without any filtering. The same applies to the results after removing the systematics, but once the time component is also subtracted, a small fraction of the red points move below the main envelope of the relation. This can be understood as follows. The time component of the model for giants tends to use a quasi-periodic kernel with a short period, which fits the oscillations, but if this period is too close to the Nyquist frequency of the data, the time component also explains some of the white (random) noise. On the other hand, in some cases the

Lomb–Scargle periodogram criterion for using a quasi-periodic GP is not met, and the variability of the giants is not modelled at all, so there is still a cluster of red points above the main envelope in the final column of Fig. 6. Better performance for pulsating stars could be obtained by modelling each oscillation frequency using, e.g. a cosine kernel GP (see e.g. Brewer & Stello 2009), but this would require an object-by-object treatment, which is beyond the scope of this paper.

3.2 Comparison to other pipelines

A number of other teams have developed their own pipelines to extract and detrend *K2* light curves (see e.g. Vanderburg 2014; Vanderburg & Johnson 2014; Armstrong et al. 2015, 2016; Huang et al. 2015; Lund et al. 2015). It would not be practical to include here a detailed comparison to all of these – added to which, not all of the resulting light curves are public, or available for the campaigns which we have processed with κ_{2SC} . We therefore focus our comparison on the κ_{2SFF} (Vanderburg 2014; Vanderburg & Johnson 2014) and $\kappa_{2VARCAT}$ (Armstrong et al. 2015, 2016) pipelines, as the light curves produced by these pipelines are available at the time of writing, as ‘High-Level Science Products’ (HLSPs) on the MAST *K2* archive, for Campaigns 1 to 4.

Applying our κ_{2SC} pipeline to the *SAP* or *PDC* light curves systematically gives better results than the $\kappa_{2VARCAT}$ pipeline. This is illustrated for two example cases from Campaign 4, on Fig. 9. The examples include both variable and quiet stars; in either case both pipelines give broadly similar results, but the κ_{2SC} light curves are considerably less noisy. In the first example, the $\kappa_{2VARCAT}$ also contains a large discontinuity, which is not present in the κ_{2SC} version. We note that, in some cases, the raw $\kappa_{2VARCAT}$ light curves are less noisy, and show less obvious systematics, than the *SAP* or *PDC* light curves, perhaps because the apertures used are more optimal. In those cases, applying κ_{2SC} to the raw $\kappa_{2VARCAT}$ might give even better results, but this is outside the scope of this paper.

We also compare the κ_{2SC} -detrended *PDC* light curves from Campaign 4 to the output of the widely used κ_{2SFF} pipeline. Fig. 10 shows the relative difference in CDPP between the two sets of light curves. In both cases, we computed the proxy CDPP estimates ourselves using the procedure described above, as the procedure described in Vanderburg & Johnson (2014) to estimate CDPP differs slightly from ours (specifically, it tends to give slightly lower values if any correlated noise on transit time-scales remains). Except at the end of the magnitude range, where there are too few objects to make a meaningful comparison, the relative difference in CDPP is below 15 per cent on average. The κ_{2SFF} light curves tend to have slightly lower CDPP for bright stars ($K_p \leq 14$) and slightly higher CDPP for fainter stars. (The same pattern was already seen in Paper I using data from the *K2* engineering test). However, across the entire magnitude range, there are cases where the CDPP values are very different in either direction. This implies that it may be useful to compare the κ_{2SFF} and κ_{2SC} -corrected light curves for any individual object, before making a decision on which to use. Once again, even better results may be obtained by applying κ_{2SC} to the raw κ_{2SFF} light curves, but this is not possible at present, as the publicly available κ_{2SFF} light-curve files do not include 2D position information.

4 TRANSIT INJECTION TESTS

We tested how the κ_{2SC} detrending affects transit searches by carrying out injection tests using 7000 randomly selected Campaign 5

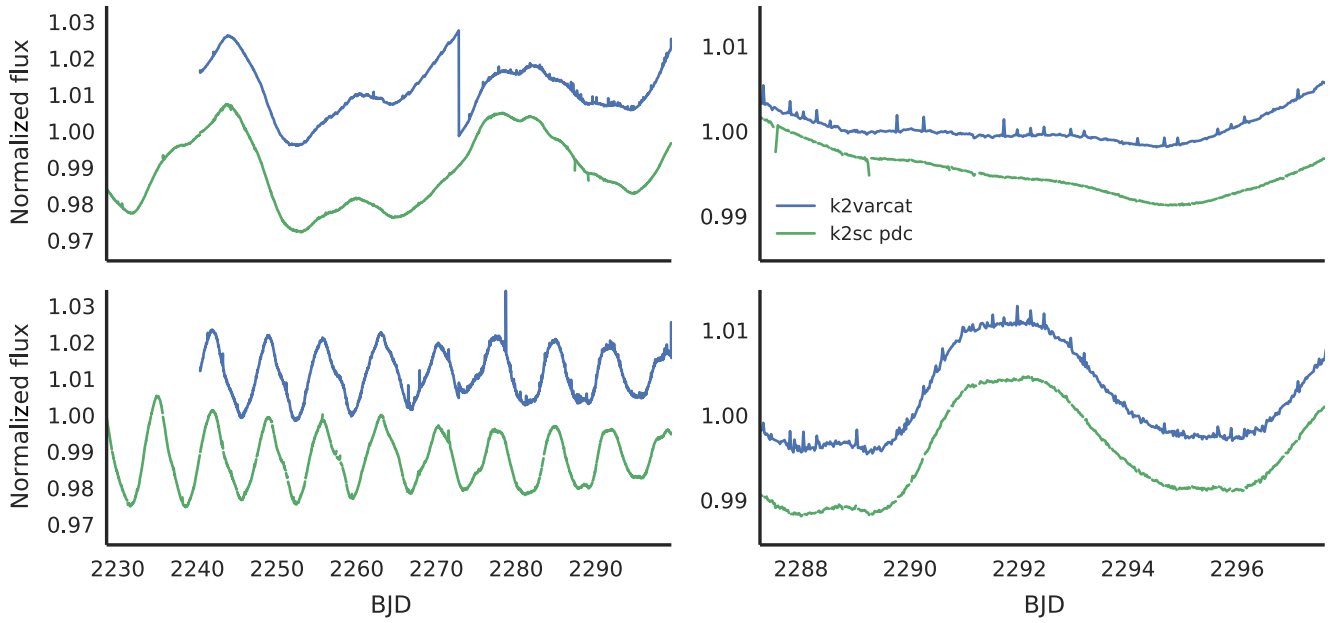


Figure 9. Comparison between the $\kappa 2_{\text{VARCAT}}$ light curves (blue) and the $\kappa 2_{\text{SC}}$ -detrended SAP light curves (green, vertically offset for clarity) for four Campaign 4 objects. The full Campaign 4 duration is shown on the left, and a zoomed subset on the right.

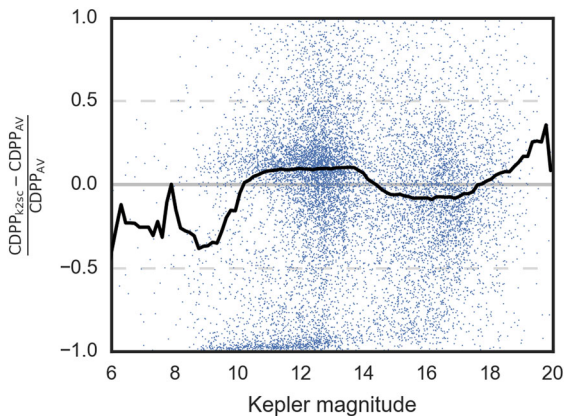


Figure 10. The relative CDPP difference between the $\kappa 2_{\text{SC}}$ and $\kappa 2_{\text{SC}}$ -corrected PDC light curves for Campaign 4. The blue dots show the individual differences, and the black line the running median with a width of 0.5 mag. The clump of points with relative CDPP difference close to 0 are red giants (see discussion in Section 3.1).

stars. In each case, we simulated a transit signal using PyTransit⁵ (Parviainen 2015), and injected it into the PDC light curve. We then detrended the light curve using $\kappa 2_{\text{SC}}$, and ran a basic transit search using our own implementation⁶ of the box least-squares (BLS) algorithm of Kovács, Zucker & Mazeh (2002). We also ran the BLS search for the data with only partial detrending, where we removed the $\kappa 2_{\text{SC}}$ time component (variability) but did not remove the position component (systematics). This partially detrended data set mimics a normal transit search case with basic time-based detrending, and works as a baseline against which we test the effects of the full detrending. Finally, we carried out the detrending and BLS search for the same data set, but without transit signal injection to

identify any existing signals present in the original light curves. This led us to remove ~ 50 stars from the analysis because they had a strong BLS detection statistic (signal detection efficiency SDE > 10) without an injected signal.

As the PDC-MAP pipeline is not publicly available, we could not inject the simulated transit signals into the light curves before the PDC-MAP step. Therefore, the tests described in this section do not account for the effects of the PDC-MAP pipeline on transit signals, but only for the effect of our $\kappa 2_{\text{SC}}$ pipeline on such signals. However, we note the PDC-MAP pipeline is designed to preserve transit signals, if necessary at the expense of other astrophysical signals. For a detailed investigation of the effect of the PDC-MAP pipeline on simulated transit signals, see Christiansen et al. (2013, 2015).

The transit signals were simulated with a sampling of 3 min, before integrating them to the $K2$ cadence of 30 min. The parameters of the simulated transits were drawn from the following distributions:

- (i) stellar density ρ_* : Gaussian with mean 1.7 and standard deviation $\sigma = 0.1 \text{ g cm}^{-3}$;
- (ii) orbital period P : uniform from 0.75 to 40 d;
- (iii) planet-to-star radius ratio k : half-Gaussian with mean 0 and $\sigma = 0.75$ (positive half only).

The impact parameter and orbital eccentricity were set to zero for simplicity, and we fixed the quadratic limb-darkening law coefficients to 0.4 and 0.1.

The BLS search was carried out over a uniform grid of 5000 frequencies from 0.025 to 1.33 d^{-1} (0.75–40 d) with q_{min} of 0.001, q_{max} of 0.2 (where q_{min} and q_{max} are the minimum and maximum transit durations in units of the orbital period, respectively) and per-frequency binning of 500, and the signal giving the highest SDE was recorded as the detected signal. The injected signal was considered as correctly identified if the recovered period was within 2 per cent of the injected period. The BLS algorithm computes a periodogram of signal-to-noise ratios (SNRs) of the best-fitting transit at each trial period. The detection statistic used is then the SDE, which is

⁵ Available at <https://github.com/hpparvi/PyTransit>

⁶ Available at <https://github.com/hpparvi>




Figure 11. BLS signal detection efficiency (SDE) for the identified and unidentified injected signals for original PDC data (upper panel) and the $\kappa 2sc$ -detrended PDC data (bottom panel).

the ratio of the highest peak in the SNR periodogram to its standard deviation.

Since all the light curves contained an injected signal, but none of them had strong BLS signal before injection, the cases where the recovered signal was not the injected one (which we refer to as ‘unidentified signals’) are cases where another signal was detected instead, i.e. false alarms. Fig. 11 shows the BLS SDE distributions for the correctly identified and unidentified signals for the original and $\kappa 2sc$ -detrended PDC data, and Fig. 12 shows the same as a function of the injected orbital period.

Fig. 13 shows the False Identification Probability (FIP) as a function of radius ratio and orbital period for original and detrended PDC data, as well as the FIP marginalized over the injected period. The FIP is defined as the fraction of the detected signals which are not correctly identified. The upper row shows the FIP for the whole data set, without imposing any detection threshold. The lower row shows a more realistic transit search case where we used a detection threshold of $SDE \geq 8$. (In an even more realistic case, one might use a variable SDE cut as a function of identified radius ratio and period.) The $\kappa 2sc$ detrending improves our capability to correctly identify the injected transit signals over the whole (p, k) -space independent of SDE value, as shown from the first row. More importantly, when applying an SDE cut, we see a significant difference in our ability to recover injected signals, particularly those corresponding to small planets (radius ratio smaller than 0.05).

5 PUBLICLY AVAILABLE CODE AND DATA

5.1 The $\kappa 2sc$ PACKAGE

$\kappa 2sc$ is distributed as a PYTHON package containing the main $\kappa 2sc$ executable and a set of PYTHON modules. The executable can be used to detrend K2 light curves from several sources, and offers an automatic MPI parallelization for detrending large data sets using




Figure 12. BLS signal detection efficiency (SDE) for the identified and unidentified signals as a function of period of the injected signal for the original (upper panel) and $\kappa 2sc$ -detrended (lower panel) PDC data. Blue dots mark all the injected signals, and black dots the correctly identified signals.

computing clusters. The $\kappa 2sc$ PYTHON modules offer a lower level access to the code, which is useful when, for example, using a custom K2 photometry pipeline and the light-curve format is not supported by the executable, or when the user wants to experiment with new GP kernels.

The package is open source with a GPLv3 license, and is available from GitHub

<https://github.com/OxES/k2sc>

and from the PYTHON Package Index. The installation follows normal PYTHON package installation steps, and we have aimed to keep the external dependencies in minimum (SCIPY, NUMPY, GEORGE, and PYFITs⁷). The code has been written keeping an eye on extendability, and adding new GP kernels and importers for data from custom photometry pipelines is easy.

The detrending time for a single light curve on a modern computer is dominated by the given number of DE iterations to carry out before the local hyperparameter optimization. Generally, the results do not improve significantly after 1–2 min of hyperparameter optimization, which leads to a total detrending time of ~ 2 min per light curve.

5.2 Processed light curves from K2 Campaigns 3 to 5

We have processed all the light curves generated using the KSOC pipeline and available at the MAST. At the time of writing, this includes the long-cadence light curves from Campaigns 3 to 5. These $\kappa 2sc$ -processed LCs are available from MAST as a K2 HLSP. For more information on how to search for, download and use the $\kappa 2sc$ HLSPs, see archive.stsci.edu/prepds/k2sc/.

⁷ PYFITs is a product of the Space Telescope Science Institute, which is operated by AURA for NASA.




Figure 13. The false identification probability (FIP) as a function of radius ratio and orbital period for the whole data set (upper row), and for signals with BLS $SDE > 8$ (lower row). The first column shows the FIP for the original PDC data, the second column shows the FIP for the $\kappa 2sc$ -detrended PDC data, and the third column shows the FIP marginalised over the orbital period. The probability densities are calculated using Gaussian kernel density estimation. The completeness of the search can be calculated as $1 - \text{FIP}$ for $SDE > 0$.

6 CONCLUSIONS AND FUTURE WORK

We have presented a new pipeline, $\kappa 2sc$, to model and, if desired, remove instrumental systematics and astrophysical variability in light curves from space-based transit surveys. The development of such a pipeline was motivated by the strong roll-induced systematics present in *K2* data, but the same method could in principle be applied many other data sets. The $\kappa 2sc$ pipeline models the systematics as a GP depending on the star’s position on the detector, and the astrophysical variability as a separate, additive GP depending on time only. It automatically checks for periodic behaviour in the light curves and, if this is found, adapts the time component accordingly. We have applied this $\kappa 2sc$ to the publicly available long-cadence *K2* SAP and PDC light curves, and evaluated the photometric precision of the resulting detrended light curves by computing a proxy estimate of their CDPP on 6-h time-scales. The CDPP for bright stars ($V \sim 12$) is within a factor of 1.5 of the original *Kepler* mission.

We compared our results to publicly available light curves produced by other pipelines, specifically $\kappa 2VARCAT$ (Armstrong et al. 2015, 2016) and $\kappa 2SFF$ (Vanderburg & Johnson 2014; Vanderburg 2014). Our light curves are consistently less noisy and contain fewer ‘glitches’ and discontinuities than the $\kappa 2VARCAT$ ones. The precision we achieve (measured in terms of 6-h CDPP) is broadly similar to that of the $\kappa 2SFF$ pipeline, but the latter gives slightly better results for bright stars ($V \leq 14$) while $\kappa 2sc$ gives slightly better results for fainter stars. We speculate that this may be due to differences between the apertures used by the KSOC pipeline (which produces the SAP and PDC light curves, to which we applied $\kappa 2sc$) and the $\kappa 2SFF$ pipeline. An important advantage of $\kappa 2sc$ over $\kappa 2SFF$ is that it is more robust to astrophysical variability. Specifically, because the systematics and variability are modelled simultaneously, $\kappa 2sc$ is better able to distinguish between the two, even when the variability is strong. Furthermore, our GP-based approach to modelling variability, including where appropriate a kernel that reproduces the quasi-periodic behaviour of rotating active stars, combined with a careful treatment of outliers, enables the variability to be subtracted

to reveal even low-amplitude transits and other short-lived events. While $\kappa 2sc$ is designed to be general-purpose, we therefore recommend it particularly to anyone who is interested in studying variable stars with *K2*, or in searching for transits or flares in the light curves of variable stars.

We also used signal injection and recovery tests to explore the completeness of planetary transit searches in $\kappa 2sc$ -detrended PDC light curves, for planets with radius ratio k down to 0.01 and period P up to 40 d in some cases. We demonstrated better than 50 per cent sensitivity down to $k = 0.03$ at $P = 5$ d, rising to 0.05 at $P = 40$ d (see Fig. 13, upper middle panel). This implies that *K2* should be able to detect warm Neptunes and hot super-Earths around Sun-like stars, and some Earth-sized planets in the habitable zones of M dwarfs.

While the results to date are very promising, the photometric performance might be improved further, for example, by using a different set of light curves than the SAP or PDC as input. The $\kappa 2sc$ pipeline can in principle be applied to any light curve so long as it contains time, flux and 2D position information. For example, light curves extracted with more optimal apertures might yield better end results. We note that users wishing to read a new light-curve format with $\kappa 2sc$ can readily do so by writing a small wrapper function, or are welcome to contact us for help if required. On a separate note, the $\kappa 2sc$ -processed *K2* light curves we have examined, as well as light curves produced by other pipelines, often display some long-term trends, which are common to many stars. Such trends were also common in the original *Kepler* data, they can be caused by a range of factors from aperture losses associated with the long-term drift of the telescope pointing to temperature and focus changes. The PDC pipeline was designed to remove such trends from SAP light curves, but it may function less effectively in the presence of the large amplitude systematics associated with the *K2* roll-angle variations. It might be more effective to model the latter first using $\kappa 2sc$, and then attempt to remove the common-mode trends using a principal component analysis-like method, such as e.g. SysRem

(Tamuz, Mazeh & Zucker 2005), or the slightly more sophisticated ARC (Astrophysically Robust Correction) method we developed for the original *Kepler* mission (Roberts et al. 2013).

Another possible avenue for further improvement might be improving the initial guesses for the GP hyperparameters, for example, incorporating information about the dependence of the position-dependent systematics component of our model on an object's location in the *Kepler* FOV.

For users interested in a particular kind of variable stars, better results might be achieved on an object-by-object basis by using a specialized GP covariance function to model the time component in the $\kappa 2sc$ framework. For example, for pulsating stars, one might use a sum of decaying cosine kernels (as employed by Brewer & Stello 2009), with priors on the pulsating frequencies derived from a preliminary analysis of the light-curve power spectrum. Once again, adding custom-built GP kernels to $\kappa 2sc$ should be relatively easy, and prospective users are encouraged to try it themselves and/or contact us for help.

Finally, we conclude by noting that the methodology we have developed here is readily applicable to other data sets, potentially including ground-based transit surveys such as the Next Generation Transit Survey, as well as the upcoming Transiting Exoplanet Survey Satellite (*TESS*) and the future *PLATO* mission. Since these instruments may suffer from different kinds of systematics, it might be necessary to modify the systematics component of the model, potentially including other input variables (e.g. detector temperature, seeing, airmass, and so on...) but the overall framework should still apply.

ACKNOWLEDGEMENTS

We wish to thank Andrew Vanderburg, Daniel Foreman-Mackey, Ann Marie Cody and Tom Barclay for useful discussions and feedback. This work was made possible by financial support from the Leverhulme Trust (RPG-2012-661), the UK Science and Technology Facilities Council (ST/K00106X/1) and the Clarendon Fund. The data presented in this paper were obtained from the 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.

REFERENCES

Aigrain S., Pont F., Zucker S., 2012, *MNRAS*, 419, 3147
 Aigrain S., Hodgkin S. T., Irwin M. J., Lewis J. R., Roberts S. J., 2015, *MNRAS*, 447, 2880 (Paper I)

Ambikasaran S., Foreman-Mackey D., Greengard L., Hogg D. W., O'Neil M., 2014, preprint (arXiv: e-prints)
 Armstrong D. J., Osborn H. P., Brown D. J. A., Kirk J., Lam K. W. F., Pollacco D. L., Spake J., Walker S. R., 2014, preprint (arXiv: e-prints)
 Armstrong D. J. et al., 2015, *A&A*, 579, A19
 Armstrong D. J. et al., 2016, *MNRAS*, 456, 2260
 Brewer B. J., Stello D., 2009, *MNRAS*, 395, 2226
 Buysschaert B. et al., 2015, *MNRAS*, 453, 89
 Christiansen J. L. et al., 2012, *PASP*, 124, 1279
 Christiansen J. L. et al., 2013, *ApJS*, 207, 35
 Christiansen J. L. et al., 2015, *ApJ*, 810, 95
 Gibson N. P., Aigrain S., Roberts S., Evans T. M., Osborne M., Pont F., 2012, *MNRAS*, 419, 2683
 Gilliland R. L. et al., 2011, *ApJS*, 197, 6
 Howell S. B. et al., 2014, *PASP*, 126, 398
 Huang C. X., Penev K., Hartman J. D., Bakos G. Á., Bhatti W., Domsa I., de Val-Borro M., 2015, *MNRAS*, 454, 4159
 Huber D., Bryson S. T., Haas M. R., Barclay T., Howell S. B., Sharma S., Stello D., Thompson S. E., 2016, *ApJS*, 211, 2
 Jenkins J. M. et al., 2010, *ApJ*, 713, L120
 Kovács G., Zucker S., Mazeh T., 2002, *A&A*, 391, 369
 Libralato M., Bedin L. R., Nardiello D., Piotto G., 2016, *MNRAS*, 456, 1137
 Lomb N. R., 1976, *Ap&SS*, 39, 447
 Lund M. N., Handberg R., Davies G. R., Chaplin W. J., Jones C. D., 2015, *ApJ*, 806, 30
 Parviainen H., 2015, *MNRAS*, 450, 3233
 Powell M. J. D., 1964, *Comput. J.*, 7, 155
 Press W., Teukolsky S., Vetterling W., Flannery B., 2007, *Numerical Recipes: The Art of Scientific Computing*, 3rd edn. Cambridge Univ. Press, Cambridge, p. 1256
 Rasmussen C. E., Williams C. K. I., 2006, *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA
 Roberts S., McQuillan A., Reece S., Aigrain S., 2013, *MNRAS*, 435, 3639
 Smith J. C. et al., 2012, *PASP*, 124, 1000
 Storn R., Price K., 1997, *J. Glob. Optim.*, 11, 341
 Stumpe M. C. et al., 2012, *PASP*, 124, 985
 Stumpe M. C., Smith J. C., Catanzarite J. H., Van Cleve J. E., Jenkins J. M., Twicken J. D., Girouard F. R., 2014, *PASP*, 126, 100
 Tamuz O., Mazeh T., Zucker S., 2005, *MNRAS*, 356, 1466
 Tenenbaum P. et al., 2012, *ApJS*, 199, 24
 Van Cleve J. E. et al., 2015, *PASP*, in press
 Vanderburg A., 2014, preprint (arXiv: e-prints)
 Vanderburg A., Johnson J. A., 2014, *PASP*, 126, 948

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.