

# ESTIMATING A COMMON PERIOD FOR A SET OF IRREGULARLY SAMPLED FUNCTIONS WITH APPLICATIONS TO PERIODIC VARIABLE STAR DATA<sup>1</sup>

BY JAMES P. LONG, ERIC C. CHI AND RICHARD G. BARANIUK

*Texas A&M University, North Carolina State University and Rice University*

We consider the problem of estimating a common period for a set of functions sampled at irregular intervals. The motivating problem arises in astronomy, where the functions represent a star's observed brightness over time through different photometric filters. While current methods perform well when the brightness is sampled densely enough in at least one filter, they break down when no brightness function is densely sampled. In this paper we introduce two new methods for period estimation in this important latter case. The first, multiband generalized Lomb–Scargle (MGLS), extends the frequently used Lomb–Scargle method to naively combine information across filters. The second, penalized generalized Lomb–Scargle (PGLS), builds on MGLS by more intelligently borrowing strength across filters. Specifically, we incorporate constraints on the phases and amplitudes across the different functions using a nonconvex penalized likelihood function. We develop a fast algorithm to optimize the penalized likelihood that combines block coordinate descent with the majorization–minimization (MM) principle. We test and validate our methods on synthetic and real astronomy data. Both PGLS and MGLS improve period estimation accuracy over current methods based on using a single function; moreover, PGLS outperforms MGLS and other leading methods when the functions are sparsely sampled.

**1. Introduction.** Periodic variable stars play an important role in several areas of modern astronomy, including extragalactic distance determinations and estimation of the Hubble constant [Riess et al. (2011), Shappee and Stanek (2011)]. To effectively use periodic variables, astronomers need accurate period estimates. For instance, extragalactic distance determination is performed by regressing the luminosity of a set of periodic variables stars on their log period. Incorrect period estimates introduce error into the regression parameter estimates, which in turn introduce error into the galactic distance estimate.

Astronomers estimate periods using the *light curve* of a star. A light curve is a set of brightness measurements of a star taken over time. Many astronomical surveys measure the brightness of stars in several photometric filters, or bands.

---

Received December 2014; revised July 2015.

<sup>1</sup>Supported in part by the Grants NSF CCF-1117939, ONR N00014-11-1-0714, ARO MURI W911NF-09-1-0383 and CIA Postdoctoral Fellowship #2012-12062800003.

*Key words and phrases.* Astrostatistics, penalized likelihood, period estimation, functional data, MM algorithm, block coordinate descent.

Multiband data is useful because differences in brightness across bands is a strong indicator of stellar class. Figure 1(a) displays the light curve of the periodic variable OGLE-LMC-T2CEP-041 observed in the I band ( $\times$ ) and V band ( $\circ$ ) by the Optical Gravitational Lensing Experiment (OGLE) [Udalski et al. (2008)]. This star has been observed 702 times in the I band and 76 times in the V band over the course of roughly 4000 days. The intervals between brightness measurements are irregular, and the star is observed at different times in the different bands. This is typical for light curves. Many stars are behind the sun for part of the year, leading to months-long gaps between observations. Additionally, weather can disrupt planned observation times. The vertical bars around each point are two-standard deviation uncertainty measurements on brightness. There are several sources of uncertainty in each magnitude measurement. These include variation in the number of photons captured by the detector in an interval of given length (even for a nominally constant star), background light (i.e., photons) captured by the detector, which must be estimated and removed to obtain an unbiased estimate of a star's brightness, and detector noise. Astronomers attempt to account for all of these sources and typically report a standard deviation along with a brightness measurement. The size of the error bars relative to the amount of variation in brightness demonstrates that OGLE-LMC-T2CEP-041 is a variable star. Stellar brightness is typically measured on the magnitude scale which is inversely proportional to log flux. Thus, brighter measurements are lower magnitude. In Figure 1(a) brighter measurements, and hence lower magnitudes, are higher on the plot, hence the inverted y-axis. OGLE-LMC-T2CEP-041 is a periodic variable star because its brightness variation follows a periodic pattern. Using the data in Figure 1(a), a period of approximately 2.48 days can be determined with a period estimation algorithm such as generalized Lomb-Scargle (see Section 3 for a description). The pattern of variation in OGLE-LMC-T2CEP-041 can be observed by plotting the brightness of the star versus phase (time modulo period). This is known as the *folded* or *phased light curve*. Figure 1(b) displays the folded light curve in each band. Accurate estimation of OGLE-LMC-T2CEP-041's period is fairly easy because the star has been observed hundreds of times in the I band. Surveys do not always collect enough measurements per band to make period determination easy. This is because observing time is limited and other scientific goals (galaxy evolution, weak gravitational lensing, large scale structure of the universe) may require observing a large area of the sky, thus limiting the number of times a single object can be repeatedly observed. For example, in Stripe 82, the Sloan Digital Sky Survey I (SDSS-I) collected light curves for  $\sim 700,000$  stars in five bands with a median of 10 observations per band. Sesar et al. (2007) identified several hundred of these stars as candidate periodic variables belonging to the class RR Lyrae but could not estimate periods due to the lack of available methodology for estimating periods with poorly sampled light curves. Later, the Sloan Digital Sky Survey II (SDSS-II) roughly tripled the number of observations per light curve to a median of 30 observations per band. In a follow-up study, Sesar et al. (2010) used this

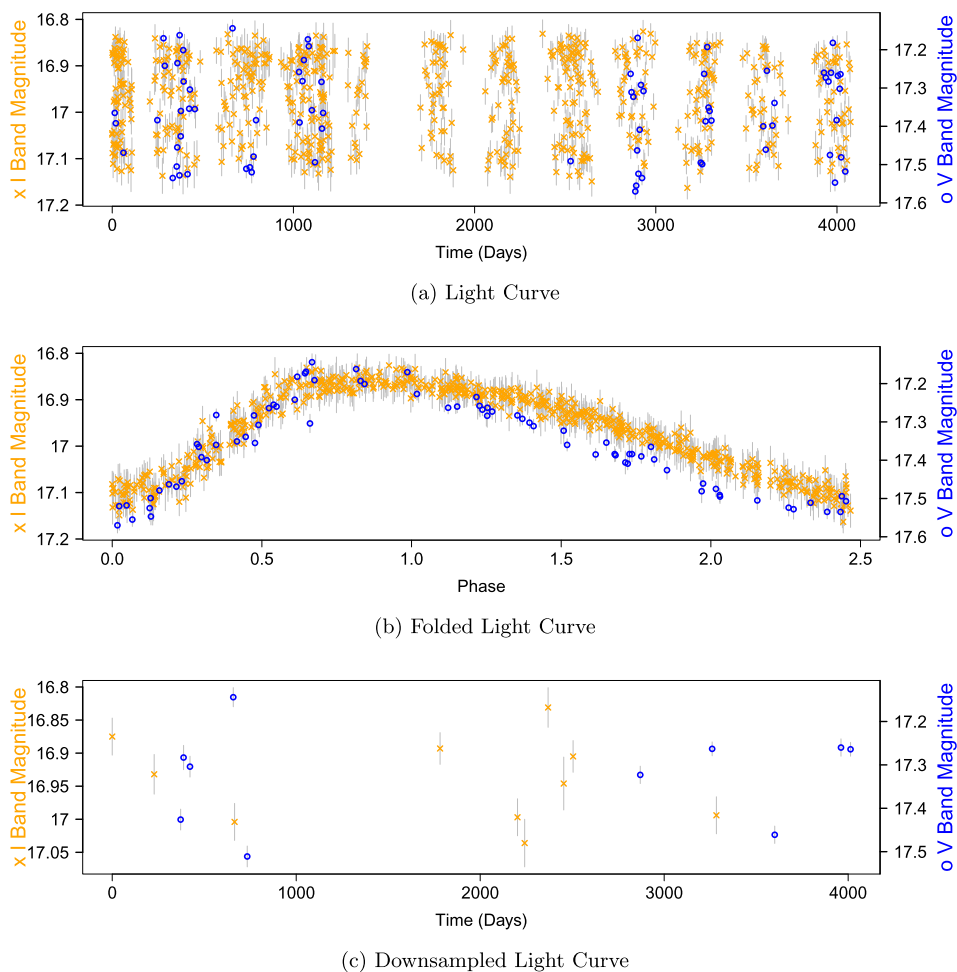


FIG. 1. (a) The light curve of the star OGLE-LMC-T2CEP-041 in the I band ( $\times$ ) and V band ( $\circ$ ). Given these measurements, astronomers seek to estimate the common period of the I and V bands. Magnitude is inversely proportional to brightness, so smaller magnitudes, representing brighter observations, are plotted higher on the plot. (b) If one is able to estimate the period ( $\approx 2.48$ ) days from the data in (a), then the pattern of brightness variation becomes apparent when brightness (i.e., magnitude) is plotted against the time modulo period. This is known as the folded light curve. Note, the phase of maximum brightness in the I band and V band are similar. (c) The same light curve as in (a) downsampled to 10 observations in the I and V bands. For data of the quality in (a), period estimation is easy because the star has been observed many times in the I and V bands. Period estimation is more difficult for data of the quality in (c), and common period estimation methods may fail. In this work we propose using correlations between the phases of maximum brightness in different bands in order to improve period estimation.

expanded set of data to estimate periods. More recently, the PanSTARRS1 survey has collected five band variable star data with a median of 4 observations per band [Schlafly et al. (2012)]. In this article, we develop methodology for estimating periods with this quality of data.

To give an idea of the challenge, Figure 1(c) shows the light curve from Figure 1(a) downsampled to 10 observations per band. In addition to enabling period recovery for already collected poorly sampled light curves, the existence of period estimation algorithm for poorly sampled data has implications for upcoming astronomical survey design. For future surveys, astronomers will make decisions such as follows: take  $n$  images per band of one field (area of the sky) or take  $n/2$  images per band of two fields. With accurate algorithms for period estimation with  $n/2$  observations per band, the second observing strategy will enable astronomers to collect twice as many objects and still determine periods.

Period estimation is challenging when stars are poorly observed in multiple bands because one must model brightness variation for several functions, thus using many degrees of freedom. For example, consider a model with  $r$  parameters to describe the shape of the light curve in each band. If one collects data in 5 bands with 10 observations per band, then the model will have  $5r + 1$  parameters (1 for the period), constrained by a total of 50 observations. In contrast with 50 observations all taken in a single band, there are only  $r + 1$  parameters to fit. Light curve shapes across different bands, however, are not independent. For example, periodic variables typically reach their peak brightness at similar points in phase space in each band [see Figure 1(b)]. Enforcing such physical constraints in models may improve the accuracy of period estimation procedures by reducing the effective degrees of freedom that must be used to fit the curves.

We now present an illustrative example that confirms this intuition and motivates our strategy for borrowing strength across multiple photometric bands. As a test, we use a sample of well-observed ( $\geq 50$  measurements/band) periodic variable stars observed in the I and V bands collected by the OGLE survey [Udalski et al. (2008)]. We estimate their periods using a simple multiband extension of generalized Lomb–Scargle (MGLS). MGLS models the brightness variation as sinusoidal and finds the best fitting period (see Section 2.3 for details). These periods are nearly correct for every star because the light curves are well observed. Figure 2(a) shows the MGLS phase estimates in the I and V bands. Phase is the location (here measured on a  $[-\pi, \pi)$  scale) in the phased light curve where the brightness reaches a maximum.<sup>2</sup> From the plot, it is evident that there is a strong correlation between the I and V band phases. This is also evident for OGLE-LMC-T2CEP-041 in Figure 1(b), where both I and V bands peak near phase  $-\pi/3$ . The outliers near  $(-\pi, \pi)$  and  $(\pi, -\pi)$  in Figure 2(a) are due to the fact that phases

---

<sup>2</sup>A light curve with a phase of  $-\pi$  is brightest at time modulo period equals 0. A light curve with a phase of 0 is brightest at time modulo period equals half the period.

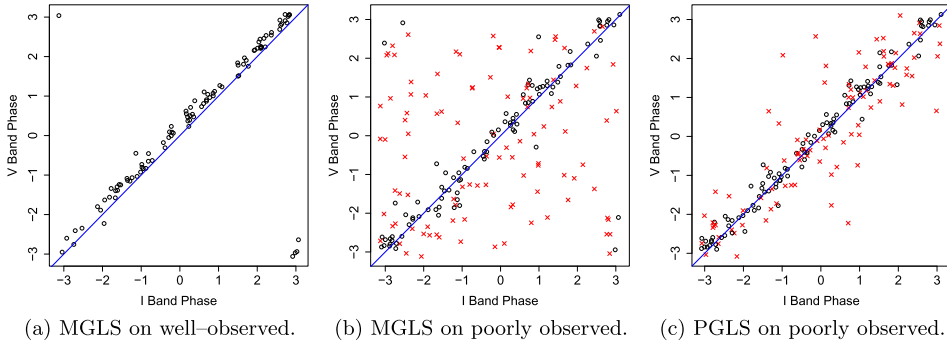


FIG. 2. (a) Phase estimates using GLS period estimates on light curves with many observations in both I and V bands. (b) When the light curves are downsampled to 10 observations per band, GLS produces phase estimates that are only weakly correlated. Furthermore, 51.5% of periods are estimated incorrectly ( $\times$ ). (c) Using the same data as in Figure 2(b), PGLS produces more correlated phase estimates. Incorrect period estimates drop to 47%.

of  $-\pi$  and  $\pi$  represent identical functions. Also note that the V band phase tends to be slightly larger than the I band phase (points tend to be above the identity line). This is due to the fact that for this type of periodic variable peak brightness in the V band typically occurs slightly before peak brightness in the I band. Now consider downsampling these light curves to 10 measurements in both the I and V bands. The data now resemble the quality, in terms of number of photometric measurements per band, of SDSS-I. We again apply the MGLS algorithm and plot the phase estimates in Figure 2(b). The points are marked by whether the period is estimated to within 1% of its true value. By this measure, 51.5% of the periods are estimated incorrectly. Note that the phase estimates do not appear strongly correlated in the two bands. The MGLS algorithm does not force the phases to be similar in different bands. An algorithm that uses the known phase correlations may be able to estimate periods more accurately. In Figure 2(c) we plot the phase estimates using a modified GLS algorithm, termed Penalized GLS (PGLS), that we develop in Section 3. PGLS enforces phases constraints across the bands using a penalized likelihood. The constraints shrink phases toward each other, resulting in tightly correlated phase estimates that are physically realistic. Incorrect period estimates have fallen to 47% from 51.5%. The remainder of this paper is structured as follows: In Section 2 we discuss additional astronomy background and existing methodology for estimating periods of variable stars. In Section 3 we introduce a penalized likelihood method for estimating periods that uses known correlations in phase and amplitude across bands. In Section 4 we develop a fast algorithm, PGLS, to maximize the likelihood by combining block coordinate descent with the majorization–minimization (MM) principle. In Section 5 we discuss selection of tuning parameters. We apply PGLS to simulated and real data in Section 6. We finish with conclusions in Section 7.

## 2. Astronomy background and existing period estimation methods.

2.1. *Photometry.* Telescopes take images of the sky using a particular filter, also known as a passband or band. From these images, stars are identified and their brightnesses estimated. Let  $m_{bi}$  be the brightness of a star observed in band  $b$  for an image taken at time  $t_{bi}$ . The standard unit of brightness in astronomy is magnitude. Magnitude is a negative log transform of luminosity, thus brighter objects have lower magnitudes. An uncertainty on  $m_{bi}$ , denoted  $\sigma_{bi}$ , is also typically recorded. Often the observed magnitude  $m_{bi}$  is modeled as being some true, unobserved magnitude plus Gaussian error with standard deviation  $\sigma_{bi}$ . Using this model, the light curve for a particular star is  $\{(t_{bi}, m_{bi}, \sigma_{bi})\}_{i=1}^{n_b}\}_{b=1}^B$ , where  $B$  is the number of bands and  $n_b$  is the number of images of the star taken in band  $b$ . Modern astronomy surveys routinely collect millions of light curves.

2.2. *Identifying variable stars.* Before period estimation is performed, astronomers typically separate the nonvariable and variable stars in a survey using variability detection methods. Simple methods involve computing the difference between the magnitudes and the mean magnitude for each band. Large absolute differences suggest that the star is a variable. This approach does not use any time information. The Welch–Stetson statistics developed in [Welch and Stetson \(1993\)](#) and [Stetson \(1996\)](#) use time correlations in residuals. These statistics are designed for multiband data where images are taken nearly simultaneously in different bands. While nonvariables play an important role in many areas of astronomy, this work is devoted to the study of period estimation for variable stars.

2.3. *Period estimation.* The set of variable stars is further sifted to identify and characterize the periodic variables. Astronomers and statisticians have developed many methods for estimating periods. Nearly all of the existing methodology has been designed for single band data observed many times. We refer interested readers to [Reimann \(1994\)](#), Chapter 3.1, and [Graham et al. \(2013\)](#) for reviews of existing methodology. Here we describe a pair of approaches, known in astronomy as Lomb–Scargle (LS) and generalized Lomb–Scargle (GLS). The multiband period estimation method developed in this article is an extension of GLS. LS and GLS both model the magnitudes as a sinusoid plus measurement error. The period estimate is the period that maximizes the likelihood. With LS the magnitudes are scaled to mean 0, and a sinusoid without an intercept is fit to the data [[Lomb \(1976\)](#), [Scargle \(1982\)](#)]. [Zechmeister and Kürster \(2009\)](#) proposed a modification of LS, GLS, in which the magnitudes are not normalized to mean 0 and an intercept term is used in the sine fit. Specifically, let

$$m_i = \beta_0 + a \sin(\omega t_i + \rho) + \varepsilon_i,$$

where  $\beta_0$ ,  $a$ ,  $\rho$  and  $\omega$  are the magnitude offset, amplitude, phase and frequency of the star. The measurement error  $\varepsilon_i \sim N(0, \sigma_i^2)$  is assumed to be independent

across  $i$ . Although the model is nonlinear in the frequency  $\omega$ , one can nonetheless compute maximum likelihood estimates by selecting a grid  $\Omega$  of possible frequencies  $\omega$  and solving a linear regression at each frequency in the grid. We describe this procedure in Section 4.1. The sinusoidal model is an approximation, as most light curves exhibit at least some degree of nonsinusoidal variation. However, LS and GLS have proven to be effective period estimation algorithms in many settings. Reimann (1994) (Section 3.3) found that for periodic light curves with unimodal behavior over phase, LS estimated frequency as well as more sophisticated approaches, including sinusoidal models with harmonic terms, periodic cubic splines and SuperSmother.<sup>3</sup> In a study of periodic variables with a variety of light curve shapes, Dubath et al. (2011) found that GLS outperformed competing period estimation methods. Part of the success of LS and GLS is due to being more computationally efficient than many alternatives. LS and GLS have been used in many recent studies of periodic variable stars [Debusscher et al. (2009), Richards et al. (2011), Suveges et al. (2012), Watkins et al. (2009)].

2.4. *Multiband period estimation methods.* While many period estimation procedures exist for single band data, few methods have been developed for multiband light curves. With multiband data, practitioners typically use single band period estimation procedures individually on each band and then combine or choose between the single band estimates based on some criteria. For example, Watkins et al. (2009) used the LS sinusoid model run separately on  $g$  and  $r$  band light curves to select candidate periods that were then analyzed using a string-length period algorithm to determine a single best estimate. Suveges et al. (2012) applied GLS to multiband data by first combining the bands together using a robust version of principal components analysis. Their method assumes that brightness measurements in the different bands are taken at the same time, limiting its applicability. The two methods we introduce below do not require simultaneous measurements.

### 3. A penalized maximum likelihood multiband period estimator.

3.1. *Multiband Lomb Scargle.* We begin by proposing a simple generalization of the single-band GLS algorithm of Zechmeister and Kürster (2009) to multiple bands. We call this first method MGLS for multiband GLS. To the best of our knowledge, MGLS is the first work to formally define GLS for multiple bands. Throughout, scalars are denoted by lowercase letters ( $u$ ), vectors by boldface lowercase letters ( $\mathbf{u}$ ), and matrices by boldface capital letters ( $\mathbf{U}$ ). We denote the standard dot product between vectors  $\mathbf{a}$  and  $\mathbf{b}$  by  $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^T \mathbf{b}$ . We denote the vector of all ones with  $\mathbf{1}$ . Let  $D = \{(t_{bi}, m_{bi}, \sigma_{bi})\}_{i=1}^{n_b}\}_{b=1}^B$  denote the data, namely, the triples of times, magnitudes and magnitude error measurements in the  $B$  bands.

---

<sup>3</sup>See Friedman (1984) for a description of SuperSmother.

Let  $\omega$  be the frequency that is shared across all bands. Brightness at time  $t$  in band  $b$  is modeled using the sine curve

$$\mu_b(t) = a_b \sin(\omega t + \rho_b) + \beta_{0b},$$

where  $a_b$ ,  $\rho_b$  and  $\beta_{0b}$  are the amplitude, phase and magnitude offset for band  $b$ . For notational efficiency we refer to vectors of amplitudes, phases and magnitude offsets as

$$\begin{aligned} \mathbf{a} &= (a_1, \dots, a_B)^\top, \\ \boldsymbol{\rho} &= (\rho_1, \dots, \rho_B)^\top, \\ \boldsymbol{\beta}_0 &= (\beta_{01}, \dots, \beta_{0B})^\top. \end{aligned}$$

We bundle these parameters into the vector  $\boldsymbol{\theta} = (\omega, \mathbf{a}^\top, \boldsymbol{\rho}^\top, \boldsymbol{\beta}_0^\top) \in \mathbb{R}^{3B+1}$ . The observed magnitude  $m_{bi}$  is a noisy measurement of  $\mu_b$  at time  $t_{bi}$ , namely,

$$m_{bi} = \mu_b(t_{bi}) + \varepsilon_{bi},$$

where  $\varepsilon_{bi} \sim N(0, \sigma_{bi}^2)$ . Assuming mutual independence of all  $\varepsilon_{bi}$ , the likelihood function becomes

$$(3.1) \quad p(D|\boldsymbol{\theta}) = \prod_{b=1}^B \prod_{i=1}^{n_b} \frac{1}{\sqrt{2\pi} \sigma_{bi}} e^{-[m_{bi} - \mu_b(t_{bi})]^2 / (2\sigma_{bi}^2)}.$$

To ensure identifiability of this model, we require that the elements  $a_b$  be non-negative and the elements  $\rho_b$  reside in the interval  $[0, \pi)$ . To keep our notation readable, we do not explicitly write out these constraints in the rest of the paper. The negative log likelihood (NLL) of  $p(D|\boldsymbol{\theta})$  is

$$(3.2) \quad \ell(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}|D) = \frac{1}{2} \sum_{b=1}^B \sum_{i=1}^{n_b} \left( \frac{m_{bi} - A_b \sin(t_{bi}\omega + \rho_b) - \beta_{b0}}{\sigma_{bi}} \right)^2.$$

The solution to the NLL at frequency  $\omega$  is

$$\ell(\omega) = \min_{\mathbf{a}, \boldsymbol{\rho}, \boldsymbol{\beta}_0} \ell(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}),$$

and the maximum likelihood estimate for  $\omega$  is

$$(3.3) \quad \hat{\omega} = \operatorname{argmin}_{\omega \in \Omega} \ell(\omega).$$

Equation (3.3) is a straightforward generalization of the GLS algorithm to multiple bands. We discuss the computation of  $\hat{\omega}$  in Section 4.1.



3.2. *Penalty terms.* We now propose a generalization of GLS that penalizes unlikely values of  $\mathbf{a}$  and  $\boldsymbol{\rho}$ . We call this method PGLS for penalized generalized Lomb–Scargle. The penalized negative log likelihood (PNLL) is

$$(3.4) \quad f(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho} | D; \gamma_1, \gamma_2) = \ell(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho} | D) + \gamma_1 J_1(\mathbf{a}) + \gamma_2 J_2(\boldsymbol{\rho}).$$

The minimal PNLL at frequency  $\omega$  is

$$(3.5) \quad f(\omega | D; \gamma_1, \gamma_2) = \min_{\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}} f(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho} | D; \gamma_1, \gamma_2),$$

and the PGLS frequency estimate is

$$(3.6) \quad \hat{\omega}(D; \gamma_1, \gamma_2) = \operatorname{argmin}_{\omega \in \Omega} f(\omega | D; \gamma_1, \gamma_2).$$

We will often suppress dependence of  $f$  and  $\hat{\omega}$  on  $D$ ,  $\gamma_1$  and  $\gamma_2$  when these quantities are fixed. The penalties  $J_1$  and  $J_2$  [see (3.7) and (3.8)] are chosen to be large for values of  $\mathbf{a}$  and  $\boldsymbol{\rho}$  that are unlikely. We motivate the form of  $J_1$  and  $J_2$  using periodic variable star data from the OGLE survey [Udalski et al. (2008)]. We fit MGLS to 100 well-observed stars in the I and V bands for four classes of periodic variables (Type I Cepheid, Type II Cepheid, RR Lyrae AB, RR Lyrae C). Since these stars have been well observed in the I and V bands (at least 50 measurements/band), the parameter estimates should be quite accurate. In Figure 3 we plot the maximum likelihood estimates of amplitude in the I and V bands for the four classes of periodic variables. In each plot, the line is chosen to pass through (0, 0) and the mean of the amplitudes in each band. Amplitudes cluster around these




FIG. 3. *Amplitude correlations in the I and V bands for classes of periodic variable stars.*




FIG. 4. Phase correlations in the I and V bands for four classes of periodic variable stars.

lines. The slope of these lines is greater for the RR Lyrae classes than for either Cepheid class. Let  $\tilde{\mathbf{a}}$  be the vector of mean amplitudes in each band. We propose penalizing the amplitude vector with

$$(3.7) \quad J_1(\mathbf{a}) = \frac{1}{2} \mathbf{a}^T (\mathbf{I} - \tilde{\mathbf{a}} \tilde{\mathbf{a}}^T) \mathbf{a} = \frac{1}{2} \|\mathbf{a} - \tilde{\mathbf{a}}^T \mathbf{a} \tilde{\mathbf{a}}\|_2^2,$$

where  $\tilde{\mathbf{a}}^T \tilde{\mathbf{a}} = 1$ . In words,  $J_1(\mathbf{a})$  is half the squared norm of the amplitude component orthogonal to  $\tilde{\mathbf{a}}$ . In Figure 4 we plot the maximum likelihood estimates for I band and V band phases by class along with a line with slope 1 and y-intercept of 0. For each class, phase estimates usually are close to the vector  $\mathbf{1}$ , implying that the light curve reaches maximum brightness in I and V bands at approximately the same point in phase space. Since a phase of  $-\pi$  and a phase of  $\pi$  are the same, the light curves represented by points in the upper left [near  $(-\pi, \pi)$ ] and lower right [near  $(\pi, -\pi)$ ] in each plot are not actually outliers.<sup>4</sup> Based on these plots,  $\boldsymbol{\rho}$  is penalized with

$$(3.8) \quad J_2(\boldsymbol{\rho}) = \frac{1}{2} \boldsymbol{\rho}^T \left( \mathbf{I} - \frac{1}{B} \mathbf{1} \mathbf{1}^T \right) \boldsymbol{\rho} = \frac{1}{2} \left\| \boldsymbol{\rho} - \frac{\mathbf{1}^T \boldsymbol{\rho}}{\mathbf{1}^T \mathbf{1}} \mathbf{1} \right\|_2^2.$$

In words,  $J_2(\boldsymbol{\rho})$  is half the squared norm of the component of  $\boldsymbol{\rho}$  that is orthogonal to  $\mathbf{1}$ . The parameter  $\gamma_2$  controls how strongly the phase estimates are forced toward

<sup>4</sup>Note that for the Type I and Type II Cepheid classes, the V band phase is slightly larger on average than the I band phase because the peak brightness occurs earlier in V band than I band. This has been documented in the astronomy literature; see Figure 5 of Freedman (1988).

a multiple of the vector  $\mathbf{1}$ . This parameter may be set based on the type of periodic variable. For example, if a practitioner is attempting to find periods for a class of variable stars where the phases in different bands are expected to be nearly identical, then  $\gamma_2$  may be made very large. In contrast, if only a weak relationship is expected, then  $\gamma_2$  may be made much smaller. In Figure 4, there is roughly the same scatter for all classes, so a single  $\gamma_2$  may be appropriate. We discuss selection of  $\gamma_2$  further in Section 5. Note that, with  $\tilde{\mathbf{a}} = (B^{-1/2}, \dots, B^{-1/2})$ ,  $J_1$  has the same form as  $J_2$ . In the remainder of this section, we discuss the existing literature on penalties similar to  $J_1$  and  $J_2$ . To the best of our knowledge, there is no existing work on using these penalties for period estimation. An algorithm for computing  $\hat{\omega}$  in (3.6) is discussed in Section 4. Methodology for choosing  $\gamma_1$  and  $\gamma_2$  is discussed in Section 5.

3.3. *Discussion of the penalties.* Both penalties in (3.7) and (3.8) can be written generically as  $J(\mathbf{v}) = \mathbf{v}^\top \mathbf{\Lambda} \mathbf{v}$  where  $\mathbf{\Lambda}$  is positive semidefinite. When  $\mathbf{\Lambda}$  is the identity matrix, we recover the classic ridge penalty [Hoerl and Kennard (1970)]. When  $\mathbf{\Lambda}$  encodes a discrete differential operator, we recover the quadratic penalty broadly used in smoothing splines, reproducing kernel Hilbert spaces and functional data analysis [Ramsay (2004)]. In fact, such quadratic penalties have garnered attention in functional principal components analysis [Allen (2013), Allen, Grose and Taylor (2014), Huang, Shen and Buja (2009), Tian et al. (2012)] when there is a natural “adjacency” among parameters as well as a prior belief that variation among adjacent parameters should be smooth. Here we propose using the discrete first order differential operator that corresponds to taking  $\mathbf{\Lambda}$  to be the graph Laplacian [Chung (1997)] of a completely connected network of  $B$  nodes where each node corresponds to a variable (amplitude or phase) of a given band. The graph Laplacian has been previously employed to enforce smoothness of regression parameters corresponding to neighbors in a gene network [Li and Li (2008)]. It has also been applied in enforcing smoothness in spatial parameters for estimating allele frequencies across populations in neighboring geographic regions [Gelfand and Vounatsou (2003), Rañola, Novembre and Lange (2014)].

3.4. *Bayesian MAP interpretation of penalties.* The penalized likelihood estimator for  $\omega$  specified by (3.6) may be interpreted as a Bayesian MAP estimator with improper prior distributions. Let  $p(D|\theta)$  be the sinusoidal likelihood from equation (3.1) and  $p(\theta)$  be some prior. Recall that  $\theta = (\omega, \mathbf{a}^\top, \boldsymbol{\rho}^\top, \boldsymbol{\beta}_0^\top) \in \mathbb{R}^{3B+1}$ . Then the posterior is  $p(\theta|D) \propto p(D|\theta)p(\theta)$ . The maximum a posterior estimator for  $\omega$  on the space  $\Omega$  (without marginalizing across other parameters) is

$$(3.9) \quad \begin{aligned} \hat{\omega} &= \operatorname{argmax}_{\omega \in \Omega} \max_{\mathbf{a}, \boldsymbol{\rho}, \boldsymbol{\beta}_0} p(D|\theta)p(\theta) \\ &= \operatorname{argmin}_{\omega \in \Omega} \min_{\mathbf{a}, \boldsymbol{\rho}, \boldsymbol{\beta}_0} -\log p(D|\theta) - \log p(\theta). \end{aligned}$$

The  $\widehat{\omega}$  in equation (3.9) matches the PGLS estimate from Section 3.2 [equation (3.6)] if the prior satisfies

$$-\log p(\boldsymbol{\theta}) = \gamma_1 J_1(\mathbf{a}) + \gamma_2 J_2(\boldsymbol{\rho}).$$

Thus,

$$p(\boldsymbol{\theta}) = e^{-1/2\mathbf{a}^\top \gamma_1 (\mathbf{I} - \tilde{\mathbf{a}}\tilde{\mathbf{a}}^\top) \mathbf{a}} e^{-1/2\boldsymbol{\rho}^\top \gamma_2 (\mathbf{I} - (1/B)\mathbf{1}\mathbf{1}^\top) \boldsymbol{\rho}}.$$

The priors on  $\mathbf{a}$  and  $\boldsymbol{\rho}$  have the form of mean 0 Gaussian densities, but the precision matrices  $\gamma_1(\mathbf{I} - \tilde{\mathbf{a}}\tilde{\mathbf{a}}^\top)$  and  $\gamma_2(\mathbf{I} - (1/B)\mathbf{1}\mathbf{1}^\top)$  are not invertible. Further, the prior on  $\boldsymbol{\beta}_0$  is uniform on  $\mathbb{R}^B$  and thus improper. While the prior  $p(\boldsymbol{\theta})$  does not explicitly depend on  $\omega$ , the fact that the minimization is done for  $\omega \in \Omega$  puts a uniform prior on  $\omega$ .

**4. Minimizing the penalized likelihood.** In this section we discuss the computation of the PGLS frequency estimate  $\widehat{\omega}$  in (3.6). The algorithm consists of the following:

1. Choosing a grid of possible frequencies  $\Omega$ .
2. Minimizing the negative log likelihood (NLL) for every frequency in the grid.
3. Minimizing the penalized negative log likelihood (PNLL) using a block coordinate descent algorithm for a sufficient subset of the frequencies in the grid.

In Section 4.1 we discuss steps 1 and 2. The frequency that minimizes the NLL is the MGLS frequency. In Section 4.2 we show that in order to obtain the PGLS frequency estimate, it is necessary only to minimize the PNLL on a subset of  $\Omega$ . In Section 4.3 we discuss step 3.

*4.1. Choice of frequency grid and minimizing the negative log likelihood.* The grid of frequencies  $\Omega$  is typically chosen on a linear scale with endpoints representing the physical minimum and maximum frequencies possible for the periodic variables of interest. For example, RR Lyrae type stars are known to have periods ranging from 0.1 days to 1 day [Chapter 6.8 of Percy (2007)]. Therefore, the grid endpoints for  $\Omega$  will occur at  $(2\pi)/1 \approx 6.28$  and  $(2\pi)/0.1 = 62.8$ . In contrast, Cepheids generally have periods ranging from 1 day to 100 days, and thus a grid of frequencies from 0.06 to 6.28 is appropriate [Chapter 6.9 of Percy (2007)]. In periodic variable classification studies where a wide set of periodic variable classes are of interest, periods may range from a few hours to hundreds of days [Richards et al. (2011)]. The grid spacing generally depends on the length of time between the first and last observation. As a star is observed for a longer and longer time, small errors in frequency result in larger and larger shifts in phase space for the folded light curve. Richards et al. (2011) used a grid with frequencies spaced  $0.1/(\max - \min)$ , where  $\max$  and  $\min$  are the maximum time and minimum times

of all star observations. We use this spacing in all examples here. For the SDSS-II RR Lyrae light curves studied in Section 6.2, there are around 20,000 frequencies in  $\Omega$ . The MGLS frequency estimate

$$\hat{\omega} = \operatorname{argmin}_{\omega \in \Omega} \min_{\mathbf{a}, \boldsymbol{\rho}, \beta_0} \ell(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho})$$

is determined by minimizing  $\ell$  with respect to  $\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}$  at every frequency in the grid. The function  $\ell$  can be minimized with respect to  $\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}$  by performing  $B$  linear regressions. In particular, with the first equality by definition and the third equality by the angle addition formula,

$$\begin{aligned} & \min_{\mathbf{a}, \boldsymbol{\rho}, \beta_0} \ell(\omega, \boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}) \\ &= \min_{\mathbf{a}, \boldsymbol{\rho}, \beta_0} \frac{1}{2} \sum_{b=1}^B \sum_{i=1}^{n_b} \left( \frac{m_{bi} - A_b \sin(t_{bi}\omega + \rho_b) - \beta_{b0}}{\sigma_{bi}} \right)^2 \\ &= \frac{1}{2} \sum_{b=1}^B \min_{a_b, \rho_b, \beta_{0b}} \sum_{i=1}^{n_b} \left( \frac{m_{bi} - A_b \sin(t_{bi}\omega + \rho_b) - \beta_{b0}}{\sigma_{bi}} \right)^2 \\ &= \frac{1}{2} \sum_{b=1}^B \min_{a_b, \rho_b, \beta_{0b}} \sum_{i=1}^{n_b} \left( (m_{bi} - A_b \cos(\rho_b) \sin(t_{bi}\omega) \right. \\ & \quad \left. - A_b \sin(\rho_b) \cos(t_{bi}\omega) - \beta_{b0}) / \sigma_{bi} \right)^2. \end{aligned}$$

The summand is the residual sum of squares when regressing  $(m_{b1}, \dots, m_{bn_b})$  on to the predictors  $(\sin(t_{b1}\omega), \dots, \sin(t_{bn_b}\omega))$  and  $(\cos(t_{b1}\omega), \dots, \cos(t_{bn_b}\omega))$  with known observation weights  $\sigma_{b1}^{-2}, \dots, \sigma_{bn_b}^{-2}$ .

4.2. *Minimizing the PNLL for a necessary subset of  $\Omega$ .* Compared to minimizing the NLL at a fixed  $\omega$ , minimizing the PNLL is computationally expensive (see Section 4.3). Thus, we do not want to compute the PNLL at every element in a potentially large grid  $\Omega$ . Here we show that it is possible to only compute the PNLL on a subset of  $\Omega$  and still be guaranteed of finding the frequency that minimizes the PNLL. The basic idea is to use the NLL  $\ell$  as a pointwise lower bound of the PNLL  $f$  in order to quickly check whether a frequency is a potential minimizer of  $f$ . More specifically, since  $\ell(\cdot) \leq f(\cdot)$ , if for two frequencies  $\omega$  and  $\omega'$  the inequality  $\ell(\omega) > f(\omega')$  holds, then  $\omega$  cannot be the frequency that minimizes  $f$ , since  $f(\omega) \geq \ell(\omega) > f(\omega')$ . The details for an iterative procedure based on this idea are given in Algorithm 1. The following proposition shows that  $\hat{\omega}$  determined by line 11 of Algorithm 1 is the global minimizer of  $f$  across  $\Omega$ .

PROPOSITION 4.1. *Under the notation of Algorithm 1,*

$$\operatorname{argmin}_{j \in \{1, \dots, R\}} f(\omega_j) = \operatorname{argmin}_{\omega \in \Omega} f(\omega).$$

---

**Algorithm 1** Compute the penalized neg. log likelihood (PNLL) on a subset of  $\Omega$

---

- 1: Compute  $\ell(\omega) = \min_{\mathbf{a}, \rho, \beta_0} \ell(\omega, \beta_0, \mathbf{a}, \rho)$  at every  $\omega \in \Omega$   $\triangleright$  NLL Algorithm
  - 2: Let  $\omega_1, \dots, \omega_G$  be ordering such that  $\ell(\omega_k) \leq \ell(\omega_{k+1})$
  - 3:  $i \leftarrow 0$
  - 4:  $\Omega_u(0) \leftarrow \emptyset$
  - 5: **repeat**
  - 6:      $i \leftarrow i + 1$
  - 7:      $f(\omega_i) \leftarrow \min_{\mathbf{a}, \rho, \beta_0} f(\omega_i, \beta_0, \mathbf{a}, \rho)$   $\triangleright$  PNLL BCD Algorithm
  - 8:      $\Omega_u(i) = \{\omega \in \Omega : l(\omega) > f(\omega_i)\} \cup \Omega_u(i - 1)$
  - 9: **until**  $\omega_{i+1} \in \Omega_u(i)$  or  $i = G$
  - 10:  $R \leftarrow i$
  - 11:  $\hat{\omega} = \operatorname{argmin}_{j \in \{1, \dots, R\}} f(\omega_j)$
- 

**PROOF.** By line 8 of Algorithm 1, for any  $r > R$  there exists  $h \leq R$  such that  $\ell(\omega_r) > f(\omega_h)$ . Since  $f(\cdot) \geq \ell(\cdot)$  for all  $\omega$ , we have

$$f(\omega_r) \geq \ell(\omega_r) > f(\omega_h).$$

Thus,  $\omega_r$  is not the minimizer of  $f$ .  $\square$

Algorithm 1 can greatly reduce the number of frequencies for which the PNLL must be minimized. For instance, with the SDSS-II RR Lyrae of Section 6.2 we use an  $\Omega$  grid with around 20,000 frequencies for each star. However, by using Algorithm 1, typically only a few hundred of these frequencies are possible minimizers of the PNLL. Additionally, by running Algorithm 1, one obtains the  $\mathbf{a}$ ,  $\rho$  and  $\beta_0$  minimizers of the NLL at each  $\omega$ , which serve as good initialization values in the block coordinate descent algorithm for solving  $f(\omega)$ , which we turn to next.

4.3. *A block coordinate descent algorithm.* We now describe an algorithm for minimizing the PNLL, (3.4), at a particular frequency  $\omega$ . Since the frequency  $\omega$  is fixed in this section, we drop dependence on  $\omega$  in the NLL [ $\ell(\omega, \beta_0, \mathbf{a}, \rho)$  becomes  $\ell(\beta_0, \mathbf{a}, \rho)$ ] and PNLL [ $f(\omega, \beta_0, \mathbf{a}, \rho)$  becomes  $f(\beta_0, \mathbf{a}, \rho)$ ]. Let  $w_{bi} = \sigma_{bi}^{-2}$ ,  $s_{bi}(\rho) = \sin(\omega t_{bi} + \rho)$ , and  $c_{bi}(\rho) = \cos(\omega t_{bi} + \rho)$ . Thus, we will view the PNLL as the following function of  $\beta_0, \mathbf{a}$  and  $\rho$ :

$$(4.1) \quad f(\beta_0, \mathbf{a}, \rho) = \ell(\beta_0, \mathbf{a}, \rho) + \gamma_1 J(\mathbf{a}) + \gamma_2 J(\rho),$$

where

$$\ell(\beta_0, \mathbf{a}, \rho) = \frac{1}{2} \sum_{b=1}^B \sum_{i=1}^{n_b} w_{bi} [\beta_{b0} + a_b s_{bi}(\rho_b) - m_{bi}]^2,$$

$$J_1(\mathbf{a}) = \frac{1}{2} \mathbf{a}^\top (\mathbf{I} - \tilde{\mathbf{a}} \tilde{\mathbf{a}}^\top) \mathbf{a},$$

$$J_2(\rho) = \frac{1}{2} \rho^\top \left( \mathbf{I} - \frac{1}{B} \mathbf{1} \mathbf{1}^\top \right) \rho.$$

The criterion to be minimized in (4.1) is nonconvex. Consequently, finding a global optimizer may be too much to ask. Nonetheless, we can obtain a local minimizer efficiently using an inexact block coordinate descent (BCD) algorithm. In practice, a local minimizer typically suffices. Our numerical and real data examples demonstrate that the local optimizer provides better solutions than the current state of the art. We first describe the BCD algorithm at a high level before detailing the individual steps. We minimize the PNLL (4.1) in round robin fashion with respect to the three blocks of variables  $\beta_0$ ,  $\mathbf{a}$  and  $\rho$ , holding two of them fixed while minimizing with respect to the third one. At the  $k + 1$ th round of updates, we solve in sequence the three smaller optimization problems:

$$\begin{aligned} \text{Update 1: } \quad & \beta_0^{(k+1)} = \underset{\beta_0}{\operatorname{argmin}} \ell(\beta_0, \mathbf{a}^{(k)}, \rho^{(k)}), \\ \text{Update 2: } \quad & \mathbf{a}^{(k+1)} = \underset{\mathbf{a}}{\operatorname{argmin}} \ell(\beta_0^{(k+1)}, \mathbf{a}, \rho^{(k)}) + \lambda_1 J_1(\mathbf{a}), \\ \text{Update 3: } \quad & \rho^{(k+1)} = \underset{\rho}{\operatorname{argmin}} \ell(\beta_0^{(k+1)}, \mathbf{a}^{(k+1)}, \rho) + \lambda_2 J_2(\rho). \end{aligned}$$

We emphasize that we do not impose constraints on  $\mathbf{a}$  and  $\rho$  in performing individual block updates, but rather impose the constraints at the termination of the BCD algorithm. We discuss later why we can make this simplification when we derive the individual block updates. Our motivation for using BCD is that the individual updates are simple and fast to compute. We next fill in details on these updates. Derivations of the update rules are given in detail in Appendix B. To streamline the following discussion, we define a few terms more compactly. Let  $\mathbf{s}_b(\rho) = (s_{1b}(\rho), \dots, s_{n_b b}(\rho))^T$ . Define  $\mathbf{c}_b(\rho)$  analogously. Let  $\mathbf{m}_b = (m_{1b}, \dots, m_{n_b b})^T$ .

*Update 1: The intercepts  $\beta_0$ .* Fix  $\mathbf{a}$  and  $\rho$ . Let  $\beta_0^+ = \operatorname{argmin}_{\beta_0} f(\beta_0, \mathbf{a}, \rho)$  denote the update for  $\beta_0$ . The update for  $\beta_0$  separates in its  $B$  elements and is given by

$$(4.2) \quad \beta_{b0}^+ = \frac{\langle \mathbf{1}, \mathbf{W}_b[\mathbf{m}_b - a_b \mathbf{s}_b(\rho_b)] \rangle}{\langle \mathbf{1}, \mathbf{W}_b \mathbf{1} \rangle},$$

where  $\mathbf{W}_b \in \mathbb{R}^{n_b \times n_b}$  is a diagonal matrix with  $i$ th diagonal element  $w_{bi}$ .

*Update 2: The amplitudes  $\mathbf{a}$ .* Fix  $\beta_0$  and  $\rho$ . Let  $\mathbf{a}^+ = \operatorname{argmin}_{\mathbf{a}} f(\beta_0, \mathbf{a}, \rho)$  denote the update for  $\mathbf{a}$ . The update for  $\mathbf{a}$  requires solving a linear system of equations whose solution is given by

$$(4.3) \quad \mathbf{a}^+ = \mathbf{E}^{-1} \left[ \mathbf{I} + \frac{1}{1/\gamma_1 - \tilde{\mathbf{a}}^T \mathbf{E}^{-1} \tilde{\mathbf{a}}} \tilde{\mathbf{a}} \tilde{\mathbf{a}}^T \mathbf{E}^{-1} \right] \boldsymbol{\xi},$$

where  $\mathbf{E}$  is a diagonal  $B \times B$  matrix with  $b$ th diagonal element  $e_{bb} = \mathbf{s}_b(\rho_b)^\top \times \mathbf{W}_b \mathbf{s}_b(\rho_b) + \gamma_1$ ,

$$\boldsymbol{\xi} = \begin{pmatrix} \langle \mathbf{s}_1(\rho_1), \mathbf{W}_1 \boldsymbol{\mu}_1 \rangle \\ \vdots \\ \langle \mathbf{s}_B(\rho_B), \mathbf{W}_B \boldsymbol{\mu}_B \rangle \end{pmatrix},$$

and  $\boldsymbol{\mu}_b = \mathbf{m}_b - \beta_{b0} \mathbf{1}$ . We now address why we are able to drop the nonnegativity constraints on  $\mathbf{a}$ . Suppose that the mean brightness for a star at time  $t$  in some band varies according to

$$\mu(t) = a \sin(\omega t + \rho),$$

where the amplitude  $a$  is negative. Then

$$\mu(t) = -|a| \sin(\omega t + \rho) = |a| \sin(\omega t + \rho + \pi),$$

and so having a “negative amplitude”  $a$  with phase  $\rho$  is equivalent to having a positive amplitude  $|a|$  with phase  $\rho + \pi$ . Consequently, if at the termination of PGLS an amplitude is negative, we can simply shift the estimate of the corresponding phase by  $\pi$ .

*Update 3: The phases  $\boldsymbol{\rho}$  via MM.* Fix  $\boldsymbol{\beta}_0$  and  $\mathbf{a}$ . Unlike the previous two updates, which required minimizing a convex function, the update for  $\boldsymbol{\rho}$  requires minimizing a nonconvex function. Nonetheless, it can be effectively attacked using a majorization–minimization (MM) algorithm [Becker, Yang and Lange (1997), Lange, Hunter and Yang (2000)] to get an approximate solution to the problem

$$(4.4) \quad \boldsymbol{\rho}_0^{(k+1)} = \underset{\boldsymbol{\rho}}{\operatorname{argmin}} \ell(\boldsymbol{\beta}_0^{(k+1)}, \mathbf{a}^{(k+1)}, \boldsymbol{\rho}) + \lambda_2 J_2(\boldsymbol{\rho}).$$

The basic principle behind an MM algorithm is to convert a hard optimization problem into a sequence of simpler ones. The MM principle requires majorizing the objective function  $f(\mathbf{y})$  by a surrogate function  $g(\mathbf{y}|\mathbf{x})$  anchored at the current point  $\mathbf{x}$ . Majorization is a combination of a tangency condition  $g(\mathbf{x}|\mathbf{x}) = f(\mathbf{x})$  and a domination condition  $g(\mathbf{y}|\mathbf{x}) \geq f(\mathbf{y})$  for all  $\mathbf{y} \in \mathbb{R}^d$ . The associated MM algorithm is then defined by the iterates  $\mathbf{x}_{k+1} := \operatorname{argmin}_{\mathbf{y}} g(\mathbf{y}|\mathbf{x}_k)$ . Since

$$f(\mathbf{x}_{k+1}) \leq g(\mathbf{x}_{k+1}|\mathbf{x}_k) \leq g(\mathbf{x}_k|\mathbf{x}_k) = f(\mathbf{x}_k),$$

the MM iterates generate a descent algorithm driving the objective function downhill. To update  $\boldsymbol{\rho}$ , we take advantage of the fact that for each  $b$

$$f_b(\rho_b) = \frac{1}{2} \sum_{i=1}^{n_b} w_{bi} [a_b s_{bi}(\rho_b) - \mu_{bi}]^2$$

is Lipschitz differentiable. This fact can be leveraged to construct a simple convex quadratic majorization of  $f$  as a function of  $\boldsymbol{\rho}$ .



PROPOSITION 4.2. *The following function majorizes  $\sum_b f_b(\rho_b)$  at the point  $\tilde{\rho}$ :*

$$g(\rho|\tilde{\rho}) = \sum_{b=1}^B \left[ f_b(\tilde{\rho}_b) + f'_b(\tilde{\rho}_b)(\rho_b - \tilde{\rho}_b) + \frac{L_b}{2}(\rho_b - \tilde{\rho}_b)^2 \right],$$

where

$$L_b = a_b[a_b\kappa_b + \sqrt{n_b}\|\mathbf{W}_b\boldsymbol{\mu}_b\|_2] \quad \text{and} \quad \kappa_b = \mathbf{1}^\top \mathbf{W}_b \mathbf{1}.$$

The proof is given in Appendix A. The MM algorithm generates an improved estimate  $\rho^+$  of the solution to (4.4) given a previous estimate  $\tilde{\rho}$  using the following update rule:

$$(4.5) \quad \rho^+ = \underset{\rho}{\operatorname{argmin}} g(\rho|\tilde{\rho}) + \gamma_2 J_2(\rho).$$

There is always a unique minimizer  $\rho^+$  to (4.5), since  $g(\rho|\tilde{\rho})$  is strongly convex in  $\rho$ . Moreover,  $\rho^+$  can be expressed explicitly in terms of  $\tilde{\rho}$  as

$$(4.6) \quad \rho^+ = \mathbf{F}^{-1} \left[ \mathbf{I} + \frac{1}{B/\gamma_2 - \mathbf{1}^\top \mathbf{F}^{-1} \mathbf{1}} \mathbf{1} \mathbf{1}^\top \mathbf{F}^{-1} \right] \boldsymbol{\zeta}(\tilde{\rho}),$$

where  $\mathbf{F}$  is a diagonal  $B \times B$  matrix with  $b$ th diagonal element  $f_{bb} = L_b + \gamma_2$  and

$$\boldsymbol{\zeta}(\tilde{\rho}) = \begin{pmatrix} L_1 \tilde{\rho}_1 - f'_1(\tilde{\rho}_1) \\ \vdots \\ L_B \tilde{\rho}_B - f'_B(\tilde{\rho}_B) \end{pmatrix}.$$

Similar to the nonnegativity constraint on the amplitude, we “enforce” the box constraint on the phase at the termination of the BCD algorithm. Since the criterion (4.1) is periodic in  $\rho$ , we choose the solution that is within the constraint set  $[0, \pi]^B$ . Finally, we note that it is sufficient to perform a single MM update (4.6) and then set  $\tilde{\rho} = \rho^{(k)}$  and  $\rho^{(k+1)} = \rho^+$ . Applying multiple MM updates (4.6) to obtain a more precise update for  $\rho$  typically does not pay in practice, as  $\beta_0$  and  $\mathbf{a}$  may change substantially during early rounds of updates. Moreover, as discussed below, taking only a single MM update per round of block coordinate updates does not change the overall convergence of the BCD algorithm.

*Complexity and convergence of BCD.* Algorithm 2 provides pseudocode of the BCD algorithm. We update the intercepts  $\beta_0$ , amplitudes  $\mathbf{a}$  and phases  $\rho$  in round robin fashion until the relative change in the variables falls below a defined tolerance. A single round of block updates is computationally efficient. Updating  $\beta_0$  requires  $\mathcal{O}(N)$  operations where  $N = \sum_{b=1}^B n_b$ . Updating  $\mathbf{a}$  requires computing the diagonal matrix  $\mathbf{E} \in \mathbb{R}^{B \times B}$  and vector  $\boldsymbol{\xi} \in \mathbb{R}^N$ , which requires  $\mathcal{O}(N)$  operations. Solving the linear system in line 7 of Algorithm 2 requires  $\mathcal{O}(B)$  operations

**Algorithm 2** Block Coordinate Descent + MM for penalized GLS (PGLS)

---

Initialize  $\boldsymbol{\beta}^{(0)}$ ,  $\mathbf{a}^{(0)}$  and  $\boldsymbol{\rho}^{(0)}$  as the solutions to multiband GLS (MGLS).

---

```

1:  $k \leftarrow 0$ 
2: repeat
3:   for  $b = 1, \dots, B$  do ▷ Update  $\boldsymbol{\beta}_0$ 
4:      $\beta_{b0}^{(k+1)} \leftarrow \langle \mathbf{1}, \mathbf{W}_b[\mathbf{m}_b - a_b^{(k)} \mathbf{s}_b(\rho_b^{(k)})] \rangle / \langle \mathbf{1}, \mathbf{W}_b \mathbf{1} \rangle$ 
5:   end for
6:   Update  $\mathbf{E} \in \mathbb{R}^{B \times B}$  and  $\boldsymbol{\xi} \in \mathbb{R}^B$  with  $\boldsymbol{\rho}^{(k)}$  and  $\boldsymbol{\beta}_0^{(k+1)}$  ▷ Update  $\mathbf{a}$ 
7:    $\mathbf{a}^{(k+1)} \leftarrow [\mathbf{E} - \gamma_1 \tilde{\mathbf{a}} \tilde{\mathbf{a}}^\top]^{-1} \boldsymbol{\xi}$ 
8:   Update  $\mathbf{F} \in \mathbb{R}^{B \times B}$  and  $\boldsymbol{\zeta} \in \mathbb{R}^B$  with  $\mathbf{a}^{(k+1)}$ ,  $\boldsymbol{\beta}_0^{(k+1)}$ , and  $\boldsymbol{\rho}^{(k)}$  ▷ Update  $\boldsymbol{\rho}$ 
9:    $\boldsymbol{\rho}^{(k+1)} \leftarrow [\mathbf{F} - \gamma_2 \mathbf{1} \mathbf{1}^\top]^{-1} \boldsymbol{\zeta}$ 
10:   $k \leftarrow k + 1$ 
11: until convergence

```

---

according to the explicit update given in (4.3). Similarly, setting up and solving the linear system in line 9 to update  $\boldsymbol{\rho}$  requires  $\mathcal{O}(N)$  and  $\mathcal{O}(B)$  operations, respectively. Thus, the total amount of work per round of block coordinate descent is  $\mathcal{O}(N)$  or, in other words, linear in the size of the data. We end this section with the following convergence result for Algorithm 2.

**PROPOSITION 4.3.** *The iterates  $(\boldsymbol{\beta}_0^{(k)}, \mathbf{a}^{(k)}, \boldsymbol{\rho}^{(k)})$  of Algorithm 2 tend to stationary points of the PNNL at a fixed frequency  $\omega$ .*

The proof is contained in Appendix C.

**5. Selection of  $\gamma_1$  and  $\gamma_2$ .** Recall that the regularization parameter  $\gamma_1$  controls the degree of shrinkage in the amplitude vector  $\mathbf{a}$ . Small values of  $\gamma_1$  lead to less shrinkage (amplitudes far from a multiple of  $\tilde{\mathbf{a}}$ ), and large values of  $\gamma_1$  lead to more shrinkage (amplitudes close to a multiple of  $\tilde{\mathbf{a}}$ ). We propose two methods for selecting the regularization parameters. In Method 1, termed PGLS1, we set  $\gamma_1$  equal to 0 and  $\gamma_2$  to be very large. This enforces no constraint on amplitudes but forces the phases to be nearly identical in the different bands. Essentially, this approach fits a model with a total of  $2B + 2$  parameters (one phase, one period,  $B$  intercepts and  $B$  amplitudes). This model may improve performance for poorly sampled light curves because it reduces overfitting of MGLS which has a total of  $3B + 1$  parameters. A disadvantage to this approach is that it does not use known correlations in amplitude across bands. Additionally, phases are not truly identical in the different bands, so this approach introduces some model misspecification. A more nuanced approach is to choose  $\gamma_1$  and  $\gamma_2$  such that the phase and amplitude for poorly sampled light curves match what they would be if they had been well sampled. Here we propose a computationally efficient and simple data-driven

alternative that leverages the availability of historical data of well-observed periodic variables. In astronomy, there are existing well-observed light curves where methods such as MGLS correctly estimate periods, amplitudes and phases. Using this historical data, we can estimate the scatter of the amplitudes around  $\tilde{\mathbf{a}}$  and the phases around  $\mathbf{1}$ . We tune  $\gamma_1$  and  $\gamma_2$  so that the scatter for the poorly observed amplitude and phase estimates approximately matches the scatter observed in the historical data. We term this second form of regularization parameter selection PGLS2.

5.1. *Methodology for selecting  $\gamma_1$  and  $\gamma_2$  using historical data.* We now detail the process for selecting  $\gamma_1$ . Let  $\mathbf{a}'_1, \dots, \mathbf{a}'_m$  be the amplitudes from a historical set of well-observed light curves. One could use MGLS to estimate these quantities. We assume these quantities have little measurement error because the light curves have been well observed. In the simulated and real data examples of Section 6, we assume we have access to  $m = 100$  light curves. Let  $\tilde{\mathbf{a}}$  be the normalized mean of the amplitudes and define the scatter of the amplitudes around  $\tilde{\mathbf{a}}$  as

$$(5.1) \quad s_a = \text{median}_{i \in \{1, \dots, m\}} \mathbf{a}'_i{}^\top (\mathbf{I} - \tilde{\mathbf{a}}\tilde{\mathbf{a}}^\top) \mathbf{a}'_i.$$

Given a set of poorly observed light curves  $D_1, \dots, D_n$ , we expect the scatter in the amplitudes to approximately match  $s_a$ . For some trial value of the tuning parameter  $\gamma_1$ , define the amplitude fit for  $D_i$  using PGLS as

$$\mathbf{a}_i(\gamma_1) = \underset{\mathbf{a}}{\text{argmin}} \min_{\omega, \beta_0, \rho} f(\omega, \beta_0, \mathbf{a}, \rho | D_i; \gamma_1, 0)$$

and the resulting scatter in amplitudes as

$$(5.2) \quad s_a(\gamma_1) = \text{median}_{i \in \{1, \dots, n\}} \mathbf{a}_i(\gamma_1)^\top (\mathbf{I} - \tilde{\mathbf{a}}\tilde{\mathbf{a}}^\top) \mathbf{a}_i(\gamma_1).$$

As  $\gamma_1$  increases, the amplitude estimates  $\mathbf{a}_i(\gamma_1)$  are pulled toward  $\tilde{\mathbf{a}}$ , causing  $s_a(\gamma_1)$  to decrease. As  $\gamma_1$  decreases, the amplitude estimates  $\mathbf{a}_i(\gamma_1)$  are pushed away from  $\tilde{\mathbf{a}}$ , causing  $s_a(\gamma_1)$  to increase. We tune  $\gamma_1$  such that  $s_a(\gamma_1)$  is approximately equal to  $s_a$ . Since  $s_a(\gamma_1)$  is inversely proportional to  $\gamma_1$ , we perform a binary search to find the optimal value of  $\gamma_1$ , using a log linear grid to find initial lower and upper bounds. We terminate the search when an update to  $\gamma_1$  does not alter any of the period estimates more than 1%, implying that further changes to  $\gamma_1$  will not significantly alter the resulting period estimates. An analogous procedure is used for selecting  $\gamma_2$  by substituting the vector  $B^{-1/2}\mathbf{1}$  for the vector  $\tilde{\mathbf{a}}$ . Since the total number of stars for which we want to estimate periods can be large, we select 100 light curves for determining  $s_a(\gamma_1)$  in (5.2), rather than the full set of light curves for which we seek to estimate periods. A schematic of how historical and new data are used to estimate the various parameters is given in Figure 5.

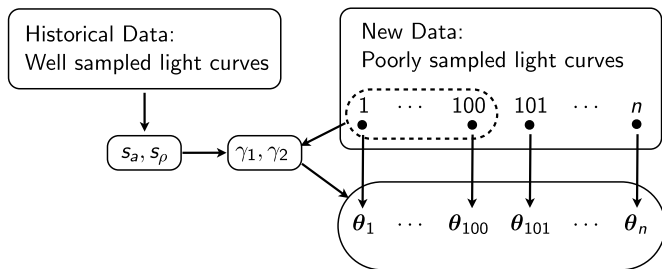


FIG. 5. *Historical periodic variable star data comprised of well sampled light curves provide a guide toward selecting tuning parameters. The scatter variables  $s_a$  and  $s_\rho$  are estimated on historical data. A subset of the new data of poorly sampled light curves, for example, the first 100 light curves, is used to select the tuning parameters  $\gamma_1$  and  $\gamma_2$ . We set  $\gamma_1$  and  $\gamma_2$  such that the scatter of amplitudes and phases approximately matches the scatter of the historical data. The  $i$ th lightcurve in the new data is used to estimate the  $i$ th set of model parameters  $\theta_i$ , namely, the  $i$ th period, amplitude, phase and intercept.*

**5.2. Choice of historical data.** In choosing  $\gamma_1$  and  $\gamma_2$  in this work, our historical data is by construction identically distributed with the poorly sampled light curves (see Section 6). In practice, for an upcoming astronomical survey, this setting could be replicated by obtaining a large number of measurements on a small fraction of periodic variable stars and using these as historical data to determine  $\tilde{\mathbf{a}}$ ,  $s_a$  and  $s_\rho$ . While this strategy would ensure the historical sample is representative of the poorly sampled light curves, it would require additional data collection. Another strategy is to use data from a past survey as historical data to determine  $\tilde{\mathbf{a}}$ ,  $s_a$  and  $s_\rho$ . This strategy will work well if light curve shape (in particular, amplitude and phase scatter) for a particular class of variable stars are similar in both the past survey and the upcoming survey. From a Bayesian perspective (see Section 3.4), using historical data which does not precisely match current data is similar to choosing a misspecified prior on  $\mathbf{a}$  and  $\boldsymbol{\rho}$ . As long as this prior is not too far from the truth, we expect that it will improve performance for period estimation.

**6. Data analysis.** We compare the performance of period estimation algorithms on simulated and real data. The four algorithms we use are GLS, MGLS, PGLS1 and PGLS2. GLS is the period found using a single band. We use the band that tends to have the largest amplitude and thus would be most likely to detect the true period. MGLS fits a separate sinusoid to each band with no penalties on amplitude and phase [equation (3.3)]. PGLS1 and PGLS2 are both forms of PGLS with different procedures for selecting  $\gamma_1$  and  $\gamma_2$  (see Section 5 for details).

**6.1. Simulations.** We simulate 500 5-band light curves from the sinusoidal likelihood model. Figure 6(a) and 6(b) show the distribution of phases and amplitudes. These amplitudes and phases are drawn from distributions meant to approximately match correlations seen in real data such as in Figures 3 and 4. The




FIG. 6. Correlations in simulated (a) phases and (b) amplitudes. These approximately match correlations observed in real data, such as in Figures 3 and 4.

periods are drawn uniformly on  $[0.2, 1.0]$  days. This represents the plausible range of periods for RR Lyrae type stars, which we study using real data in the following section. We divide the light curves into two groups: 100 historical light curves used for estimating  $\tilde{\mathbf{a}}$ ,  $s_a$  and  $s_\rho$  [see (5.1)] and 400 light curves for testing. We downsample the test light curves to 5, 10 and 15 observations per band to simulate difficult period recovery regimes. Table 1 contains the fraction of light curve periods estimated to within 1% and 5% of truth for each method for 5, 10 and 15 observations per band. PGLS2 dominates GLS, MGLS and PGLS1 in each setting. Scatterplots of true period versus period estimate for MGLS and PGLS2 for 5, 10 and 15 observations per band are contained in Figure 7. Both MGLS and PGLS2 more often underestimate than overestimate the truth (more points below identity line than above in Figure 7). This is likely caused by the fact that the sinusoid function changes more rapidly for small periods. Lower period ranges correspond to more complex models. Hastie, Tibshirani and Friedman (2009) (Section 7.8) discusses a similar issue for classification using sinusoids. In Figure 8 we plot phase estimates for MGLS and PGLS2 for 5 and 10 observation per band test sets. As

TABLE 1  
Fraction of period estimates within 1% and 5% of truth for simulation data

Obs./Band	Within 1%				Within 5%			
	GLS	MGLS	PGLS1	PGLS2	GLS	MGLS	PGLS1	PGLS2
5	0.01	0.04	0.07	0.20	0.06	0.08	0.09	0.24
10	0.04	0.40	0.50	0.80	0.07	0.42	0.52	0.80
15	0.21	0.87	0.90	0.96	0.25	0.87	0.90	0.96

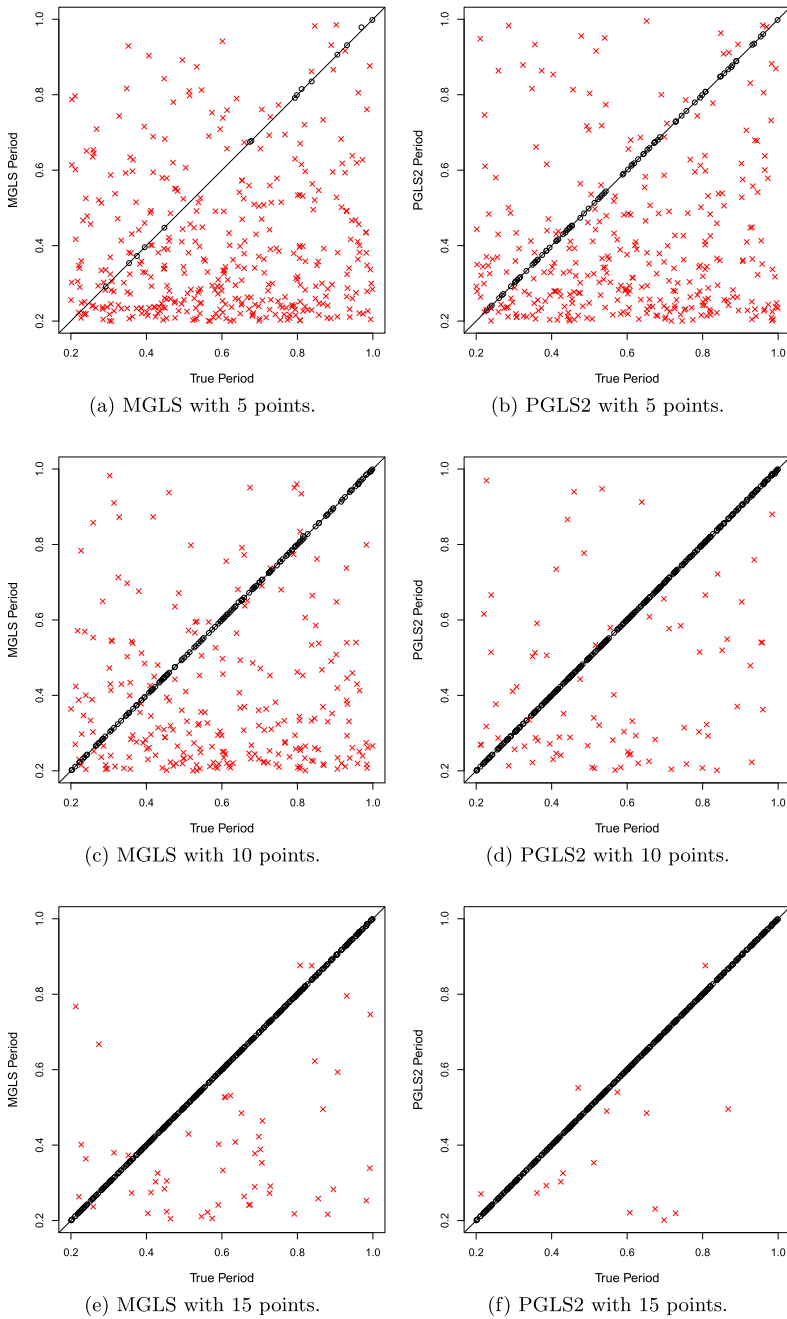


FIG. 7. Scatterplots of true period versus estimate using simulated data for MGLS and PGLS2 with 5, 10 and 15 observations per band for the test set. The MGLS estimates are in the left column and the PGLS2 estimates are in the right column. The  $\times$  are estimates further than 1% away from truth, while the  $\circ$  are estimates within 1% of truth. The PGLS2 estimates are significantly more accurate.




FIG. 8. Correlations in phase estimates for simulated data with (a) 5 observations per band and (b) 15 observations per band. The upper diagonals in each plot are PGLS2 estimates. The lower diagonals are MGLS estimates. The phases estimates using PGLS2 appear more realistic than MGLS, especially for light curves with 5 observations per band.

expected, PGLS2 phases are more concentrated around the vector  $\mathbf{1}$ . For 15 observations per band, the MGLS estimates show significant concentration around  $\mathbf{1}$ , indicative that MGLS is estimating phases correctly with 15 observations per band. Notice that PGLS2 eliminates phase estimates near  $(-\pi, \pi)$  and  $(\pi, -\pi)$ . Since these are plausible phases, this tendency could result in incorrect period estimates for some stars. A more refined penalty term for phase ( $J_2$ ) in PGLS2 could address this issue.

6.2. *SDSS stripe-82 data.* Sesar et al. (2010) identified 483 periodic variable stars of the class RR Lyrae in the Sloan Digital Sky Survey II (SDSS-II). These light curves were sampled approximately 30 times per band in five bands. Sesar et al. (2010) estimated periods for these stars using a Supersmoother routine in Reimann (1994). Visual inspection suggests that these period estimates are accurate. While Supersmoother is accurate with well-sampled light curves, with poorly sampled light curves it suffers the same problems as MGLS. This is why Sesar et al. (2007) did not attempt period estimation when SDSS-I collected a median of only 10 observations per band for these stars. Ongoing surveys, such as PanSTARRS1, currently have data roughly the quality of SDSS-I in terms of number of observations per band [Schlafly et al. (2012)]. Here we study period estimation with poorly sampled light curves by downsampling SDSS-II data and comparing GLS, MGLS, PGLS1 and PGLS2 period estimates to the Supersmoother estimates computed using the entire light curves.

We obtained 450 of 483 Sesar et al. (2010) RR Lyrae light curves from a public data repository [Ivezić et al. (2007)]. To test period estimation with poorly sampled light curves, we split the 450 light curves into 100 historical and 350 test. As

TABLE 2  
*Fraction of period estimates within 1% and 5% of truth for SDSS data*

Obs./Band	Within 1%				Within 5%			
	GLS	MGLS	PGLS1	PGLS2	GLS	MGLS	PGLS1	PGLS2
5	0.01	0.20	0.22	0.35	0.05	0.22	0.25	0.36
10	0.11	0.53	0.57	0.59	0.16	0.54	0.57	0.59
15	0.31	0.68	0.68	0.68	0.34	0.68	0.68	0.68

with the simulated data, we downsample the test light curves to 5, 10 and 15 measurements per band and compare period estimation accuracy. Table 2 shows the fraction of period estimates within 1% and 5% of their true value. The methods that use multiple bands (MGLS, PGLS1 and PGLS2) outperform GLS. PGLS2 increases the fraction of periods estimated correctly for 5 and 10 observations per band. Scatterplots of true period versus period estimate for MGLS and PGLS2 are given in Figure 9. In general, the improvement of PGLS1 and PGLS2 over MGLS appears less for the SDSS data here than the simulated data. The scatterplots in Figure 9 show errors concentrating around several nonlinear functions of the true period. This is likely due to the phenomenon of pseudo-aliasing. When observations are taken at regular intervals of 1 day, say,  $t_i = i$ , the frequencies  $\omega' = \omega + 2\pi k$  for  $k \in \mathbb{Z}$  cannot be distinguished because

$$\sin(\omega' t_i + \rho) = \sin((\omega + 2\pi k)t_i + \rho) = \sin(\omega t_i + 2\pi k t_i + \rho) = \sin(\omega t_i + \rho).$$

Converting to period scale ( $p = 2\pi/\omega$ ),  $p' = p/(1 + pk)$  are indistinguishable for any  $k \in \mathbb{Z}$ . This is known as aliasing. While the observations from SDSS are not exactly regular, most observations on the same star are taken at approximately the same time each night with stretches of missing nights due to weather and seasons where the star is behind the sun. This causes pseudo-aliasing where it is difficult (though not impossible) to distinguish between a true period  $p$  and any  $p' = p/(1 + pk)$  for  $k \in \mathbb{Z}$ . The three lines observed in the scatterplots are for  $k = -1, 1, 2$ . This issue did not arise with the simulated data because the time sampling was less regular. Period estimation methods will likely perform better on surveys with a less regular cadence. See Scargle (1982) for a discussion of pseudo-aliasing. In Figure 10 we plot phase correlations using MGLS and PGLS2 for 5 and 15 measurements per band. For both 5 and 15 observations per band, PGLS2 produces realistic phase estimates concentrated around the identity line. With 15 observations per band, the concentration exhibited by MGLS and PGLS2 is quite close. This is evidence that the penalty terms  $J_1$  and  $J_2$  are working well.

**7. Conclusions.** Accurate multiband period estimation is important for several problems in modern astronomy. Current methods, however, fail when the light curves are poorly sampled in all measurement bands. To address this deficiency, we introduced two new methods, MGLS and PGLS. Both methods generalize a well-



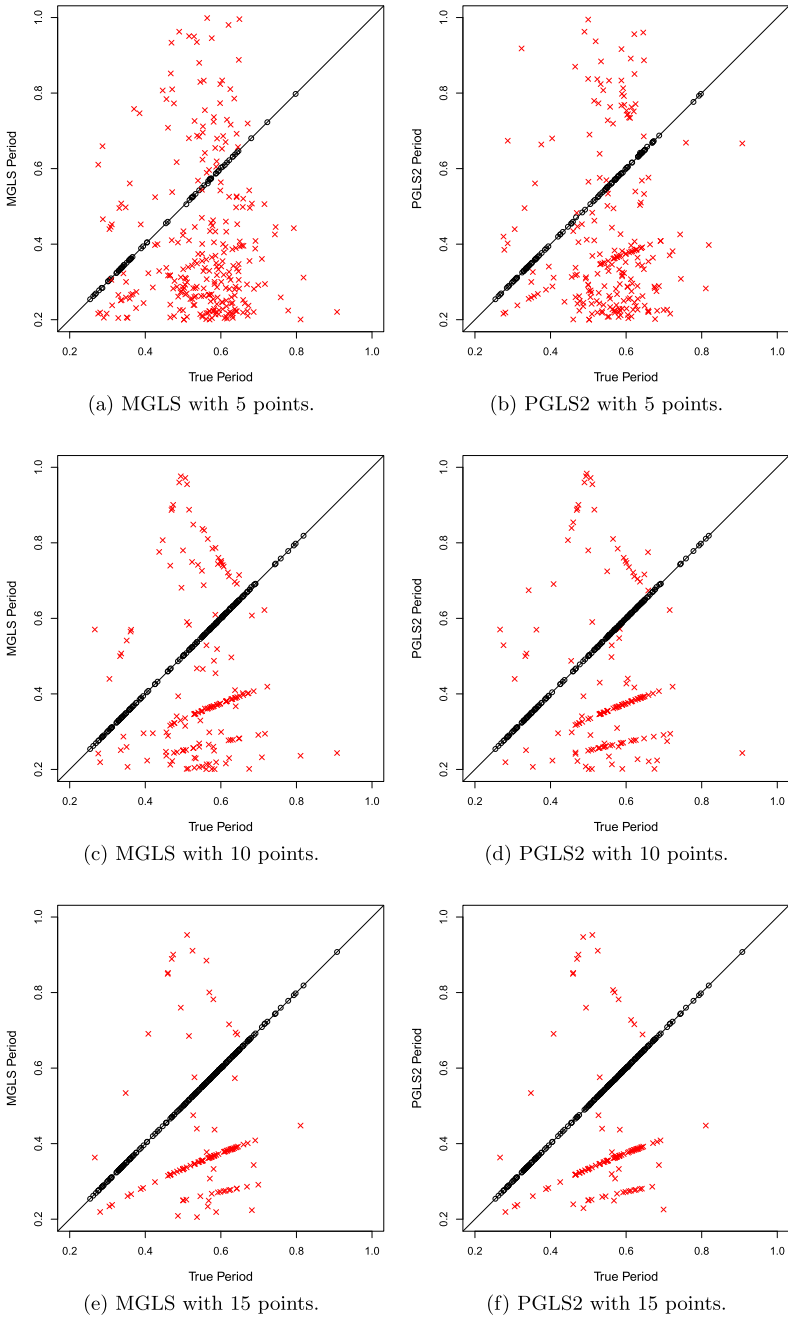


FIG. 9. Scatterplots of true period versus estimate for MGLS and PGLS2 for SDSS with 5, 10 and 15 observations per band for the test set.




FIG. 10. Correlations in SDSS phase estimates with (a) 5 observations per band and (b) 15 observations per band for the SDSS-II. The results are similar to the simulated data in Figure 8. In particular, MGLS period estimates with 5 observations per band show little clustering around 1. The PGLS2 estimates are more physically realistic.

known approach to period estimation in astronomy. However, PGLS is the first multiband period estimation algorithm for variable stars that uses known correlations in amplitude and phase across bands to improve accuracy. With simulations and real data, PGLS outperforms MGLS and GLS. Computing the PGLS estimate requires minimizing a more complicated function than computing the MGLS estimate. Nonetheless, the PGLS estimate can be rapidly obtained using our inexact BCD algorithm, which requires computational effort that scales linearly with the size of the data. We also showed that even faster PGLS estimates can be obtained by first estimating the computationally cheaper MGLS estimator at different candidate periods and using these estimates to reduce drastically the number of candidate periods on which to compute the PGLS estimate.

In this work we estimated population level uncertainty on the period estimates, that is, the fraction of light curves for which the period estimate is within a certain fraction of the truth. A more difficult problem is to estimate for a given light curve the uncertainty in the period estimate determined by MGLS or PGLS. Some progress has been made on this problem for LS and GLS. For LS, Lomb (1976) found the distribution of a test statistic for testing the hypotheses of a star having constant brightness versus the star have brightness with frequency  $\omega$ . Reimann (1994) showed that, under certain conditions on the sampling times, the LS period estimate is asymptotically normal. While this result is useful when  $n$  is large, it does not appear appropriate for the setting of limited data for which PGLS is intended. Examination of Figures 7 and 9 suggests that, with few observations, the MGLS and PGLS period estimates are not normally distributed. In future work, we plan to develop measures of uncertainty for MGLS and PGLS estimates.

Although we have focused on an application in astronomy, MGLS and PGLS have potential utility whenever multiple correlated periodic signals can be jointly modeled. Such problems arise in spatio-temporal modeling, in particular, geographic epidemiology [Torabi and Rosychuk (2010)] and air quality monitoring [Keller et al. (2015)]. In both these cases, time series data are collected at fixed geographical locations and much of the variation is seasonal or periodic. Measurements are often irregular and some collection sites take more frequent and comprehensive measurements than others. A key observation is that measurements vary smoothly with respect to geographic locality. A trade-off similar to that faced in the multiband problem is typical where a greater number of collection sites can provide a wider spatial coverage with each site taking fewer temporal samples. We anticipate that adapting our proposed methods, especially PGLS, to these problems could enable more experiments to be run concurrently by exploiting the spatial correlations.

Finally, we note that while this work was under review and posted on arXiv [Long, Chi and Baraniuk (2014)], VanderPlas and Ivezić (2015) developed two algorithms for multiband period estimation. One of these methods, the “multi-phase” model, is identical to MGLS. A second model, the “shared-phase” model, uses a ridge penalty to shrink mean magnitudes in different bands toward a common value. Appendix B of VanderPlas and Ivezić (2015) contains some comparison of the approaches.

Software for implementing MGLS and PGLS can be found in the `multiband` package for R, available on CRAN. The code for reproducing the results in this paper is available by contacting the authors.

## APPENDIX A: MAJORIZATION IN PHASE UPDATE

We provide the proof of Proposition 4.2.

**PROOF OF PROPOSITION 4.2.** Since the function we wish to majorize is Lipschitz differentiable, we can apply the quadratic upper bound principle to construct a convex quadratic majorization [Böhning and Lindsay (1988)]. The gradient and Hessian of  $f_b(\rho)$  are given by

$$f'_b(\rho) = a_b \langle a_b \mathbf{s}_b(\rho) - \boldsymbol{\mu}_b, \mathbf{W}_b \mathbf{c}_b(\rho) \rangle$$

and

$$f''_b(\rho) = a_b [\langle a_b \mathbf{c}_b(\rho), \mathbf{W}_b \mathbf{c}_b(\rho) \rangle + \langle \boldsymbol{\mu}_b - a_b \mathbf{s}_b(\rho), \mathbf{W}_b \mathbf{s}_b(\rho) \rangle].$$

It is straightforward to establish the following global upper bound on the Hessian. Let  $\kappa_b = \mathbf{1}^\top \mathbf{W}_b \mathbf{1}$ ; then

$$\begin{aligned} f''_b(\rho) &\leq a_b [a_b \mathbf{c}_b(\rho)^\top \mathbf{W}_b \mathbf{c}_b(\rho) + \langle \mathbf{W}_b \boldsymbol{\mu}_b, \mathbf{s}_b(\rho) \rangle] \\ &\leq a_b [a_b \kappa_b + \langle \mathbf{W}_b \boldsymbol{\mu}_b, \mathbf{s}_b(\rho) \rangle] \\ &\leq a_b [a_b \kappa_b + \|\mathbf{W}_b \boldsymbol{\mu}_b\|_2 \|\mathbf{s}_b(\rho)\|_2] \\ &\leq a_b [a_b \kappa_b + \sqrt{n_b} \|\mathbf{W}_b \boldsymbol{\mu}_b\|_2]. \end{aligned} \tag{A.1}$$

The exact second order Taylor expansion of  $f_b(\rho_b)$  about a point  $\tilde{\rho}$  is given by

$$(A.2) \quad f_b(\rho) = f_b(\tilde{\rho}) + f'_b(\tilde{\rho})(\rho - \tilde{\rho}) + \frac{1}{2}f''_b(\rho^*)(\rho - \tilde{\rho})^2,$$

where  $\rho^* \in \alpha\rho + (1 - \alpha)\tilde{\rho}$  for some  $\alpha \in (0, 1)$ . The relations (A.1) and (A.2) together yield the desired result.  $\square$

## APPENDIX B: DERIVATION OF UPDATE RULES FOR BCD-MM ALGORITHM

We give detailed derivations of the update rules given in (4.2), (4.3) and (4.4).

**Update 1: The intercepts  $\beta_0$ .** Fix  $\mathbf{a}$  and  $\rho$ . Let  $\beta_0^+ = \operatorname{argmin}_{\beta_0} f(\beta_0, \mathbf{a}, \rho)$  denote the update for  $\beta_0$ . The update for  $\beta_0$  separates in its  $B$  elements and is given by the minimizers to  $B$  univariate convex quadratic functions

$$(B.1) \quad \beta_{b0}^+ = \operatorname{argmin}_{\beta} [\beta \mathbf{1} + a_b \mathbf{s}_b(\rho_b) - \mathbf{m}_b]^\top \mathbf{W}_b [\beta \mathbf{1} + a_b \mathbf{s}_b(\rho_b) - \mathbf{m}_b],$$

where  $\mathbf{W}_b \in \mathbb{R}^{n_b \times n_b}$  is a diagonal matrix with  $i$ th diagonal element  $w_{bi}$ . It is straightforward to show that the solution to the optimization problem (B.1) is given by

$$\beta_{b0}^+ = \frac{\langle \mathbf{1}, \mathbf{W}_b [\mathbf{m}_b - a_b \mathbf{s}_b(\rho_b)] \rangle}{\langle \mathbf{1}, \mathbf{W}_b \mathbf{1} \rangle}.$$

**Update 2: The amplitudes  $\mathbf{a}$ .** Fix  $\beta_0$  and  $\rho$ . Let  $\mathbf{a}^+ = \operatorname{argmin}_{\mathbf{a}} f(\beta_0, \mathbf{a}, \rho)$  denote the update for  $\mathbf{a}$ . To obtain the update  $\mathbf{a}^+$ , we minimize the convex quadratic function

$$(B.2) \quad \mathbf{a}^+ = \operatorname{argmin}_{\mathbf{a}} \frac{1}{2} \sum_{b=1}^B \sum_{i=1}^{n_b} w_{bi} [a_b s_{bi}(\rho_b) - \mu_{bi}]^2 + \frac{\gamma_1}{2} \mathbf{a}^\top [\mathbf{I} - \tilde{\mathbf{a}} \tilde{\mathbf{a}}^\top] \mathbf{a},$$

where  $\mu_{bi} = m_{bi} - \beta_{b0}$ . Let  $\boldsymbol{\mu}_b = \mathbf{m}_b - \beta_{b0} \mathbf{1}$ . Setting the derivative of the quadratic function in (B.2) with respect to  $a_b$  equal to zero yields the  $B$  stationary conditions

$$\mathbf{s}_b(\rho_b)^\top \mathbf{W}_b \mathbf{s}_b(\rho_b) a_b - \langle \mathbf{s}_b(\rho_b), \mathbf{W}_b \boldsymbol{\mu}_b \rangle + \gamma_1 [a_b - \langle \tilde{\mathbf{a}}, \mathbf{a} \rangle \tilde{a}_b] = 0.$$

The stationarity conditions imply that the update  $\mathbf{a}^+$  is the solution to the following linear system of equations:

$$[\mathbf{E} - \gamma_1 \tilde{\mathbf{a}} \tilde{\mathbf{a}}^\top] \mathbf{a} = \boldsymbol{\xi},$$

where  $\mathbf{E}$  is a diagonal  $B \times B$  matrix with  $b$ th diagonal element  $e_{bb} = \mathbf{s}_b(\rho_b)^\top \times \mathbf{W}_b \mathbf{s}_b(\rho_b) + \gamma_1$  and

$$\boldsymbol{\xi} = \begin{pmatrix} \langle \mathbf{s}_1(\rho_1), \mathbf{W}_1 \boldsymbol{\mu}_1 \rangle \\ \vdots \\ \langle \mathbf{s}_B(\rho_B), \mathbf{W}_B \boldsymbol{\mu}_B \rangle \end{pmatrix}.$$

Applying the matrix inversion lemma enables us to solve the system with a single matrix-vector multiply

$$\mathbf{a}^+ = \mathbf{E}^{-1} \left[ \mathbf{I} + \frac{1}{1/\gamma_1 - \tilde{\mathbf{a}}^\top \mathbf{E}^{-1} \tilde{\mathbf{a}}} \tilde{\mathbf{a}} \tilde{\mathbf{a}}^\top \mathbf{E}^{-1} \right] \boldsymbol{\xi}.$$

**Update 3: The phases  $\boldsymbol{\rho}$  via MM.** Fix  $\boldsymbol{\beta}_0$  and  $\mathbf{a}$ . The MM algorithm generates an improved estimate  $\boldsymbol{\rho}^+$  of the solution to (4.4) given a previous estimate  $\tilde{\boldsymbol{\rho}}$  using the following update rule:

$$(B.3) \quad \boldsymbol{\rho}^+ = \underset{\boldsymbol{\rho}}{\operatorname{argmin}} g(\boldsymbol{\rho} | \tilde{\boldsymbol{\rho}}) + \gamma_2 J_2(\boldsymbol{\rho}),$$

where  $g$  is the majorization given in Proposition 4.2. There is always a unique minimizer  $\boldsymbol{\rho}^+$  to (B.3), since  $g(\boldsymbol{\rho} | \tilde{\boldsymbol{\rho}})$  is strongly convex in  $\boldsymbol{\rho}$ . We now show how  $\boldsymbol{\rho}^+$  can be expressed explicitly in terms of  $\tilde{\boldsymbol{\rho}}$ . Note that

$$(B.4) \quad \frac{\partial}{\partial \rho_b} [g(\boldsymbol{\rho} | \tilde{\boldsymbol{\rho}}) + \gamma_2 J_2(\boldsymbol{\rho})] = f'_b(\tilde{\rho}_b) + L_b(\rho_b - \tilde{\rho}_b) + \gamma_2 \left[ \rho_b - \frac{1}{B} \sum_{b'=1}^B \rho_{b'} \right].$$

The stationarity conditions imply that the update  $\boldsymbol{\rho}^+$  is the solution to the linear system of equations

$$(B.5) \quad \left[ \mathbf{F} - \frac{\gamma_2}{B} \mathbf{1}\mathbf{1}^\top \right] \boldsymbol{\rho} = \boldsymbol{\zeta}(\tilde{\boldsymbol{\rho}}),$$

where  $\mathbf{F}$  is a diagonal  $B \times B$  matrix with  $b$ th diagonal element  $f_{bb} = L_b + \gamma_2$  and

$$\boldsymbol{\zeta}(\tilde{\boldsymbol{\rho}}) = \begin{pmatrix} L_1 \tilde{\rho}_1 - f'_1(\tilde{\rho}_1) \\ \vdots \\ L_B \tilde{\rho}_B - f'_B(\tilde{\rho}_B) \end{pmatrix}.$$

Again we turn to the matrix inversion lemma to solve (B.5) with a single matrix-vector multiply

$$\boldsymbol{\rho}^+ = \mathbf{F}^{-1} \left[ \mathbf{I} + \frac{1}{B/\gamma_2 - \mathbf{1}^\top \mathbf{F}^{-1} \mathbf{1}} \mathbf{1}\mathbf{1}^\top \mathbf{F}^{-1} \right] \boldsymbol{\zeta}(\tilde{\boldsymbol{\rho}}).$$

## APPENDIX C: CONVERGENCE OF THE BCD-MM ALGORITHM

We provide the proof of Proposition 4.3.

**PROOF OF PROPOSITION 4.3.** The convergence theory of monotone algorithms, like BCD and MM, hinge on the properties of the algorithm map  $\psi(\mathbf{x})$  that returns the next iterate given the last iterate. For easy reference, we state a simple version of Meyer's monotone convergence theorem [Meyer (1976)], which is instrumental in proving convergence in our setting.

**THEOREM C.1.** *Let  $f(\mathbf{x})$  be a continuous function on a domain  $S$  and  $\psi(\mathbf{x})$  be a continuous algorithm map from  $S$  into  $S$  satisfying  $f(\psi(\mathbf{x})) < f(\mathbf{x})$  for all  $\mathbf{x} \in S$  with  $\psi(\mathbf{x}) \neq \mathbf{x}$ . Suppose for some initial point  $\mathbf{x}_0$  that the set  $\mathcal{L}_f(\mathbf{x}_0) \equiv \{\mathbf{x} \in S : f(\mathbf{x}) \leq f(\mathbf{x}_0)\}$  is compact. Then (a) all cluster points are fixed points of  $\psi(\mathbf{x})$ , and (b)  $\lim_{m \rightarrow \infty} \|\mathbf{x}_{m+1} - \mathbf{x}_m\| = 0$ .*

In the context of Algorithm 2, the function  $f$  is the PNLL (4.1). Since Algorithm 2 optimizes over the triple  $\mathbf{x} = (\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho})$  for a fixed candidate frequency  $\omega$ , we take the set  $S = \mathbb{R} \times \mathbb{R} \times [-\pi/2, \pi/2]$ . We first use Theorem C.1 to establish that the iterates of the inexact BCD algorithm tend toward the fixed points of an algorithm map. We then show that the fixed points of the algorithm map correspond to the stationary points of the PNLL. In order to apply Theorem C.1, we need to (i) identify a continuous algorithm map  $\psi(\mathbf{x})$  that corresponds to Algorithm 2, (ii) check that  $f(\mathbf{x}) > f(\psi(\mathbf{x}))$  if  $\mathbf{x} \neq \psi(\mathbf{x})$ , and (iii) identify an  $\mathbf{x}_0 \in S$  so that the set  $\mathcal{L}_f(\mathbf{x}_0)$  is compact. The first step is to identify an algorithm map  $\psi(\mathbf{x})$  and check that it is continuous. We formalize how to obtain  $\mathbf{x}^+ \equiv \psi(\mathbf{x})$  from  $\mathbf{x}$  via the composition of three submaps, each of which corresponds to a block variable update

$$\begin{aligned}\psi_1(\mathbf{x}) &= \underset{\boldsymbol{\beta}_0}{\operatorname{argmin}} \ell(\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}), \\ \psi_2(\mathbf{x}) &= \underset{\mathbf{a}}{\operatorname{argmin}} \ell(\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}) + \lambda_1 J_1(\mathbf{a}), \\ \psi_3(\mathbf{x}) &= \underset{\boldsymbol{\rho}'}{\operatorname{argmin}} g(\boldsymbol{\rho}' | \boldsymbol{\rho}) + \lambda_2 J_2(\boldsymbol{\rho}').\end{aligned}$$

Then the algorithm map that corresponds to the BCD algorithm is

$$\psi(\mathbf{x}) = \begin{pmatrix} \psi_1(\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}) \\ \psi_2(\psi_1(\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}), \mathbf{a}, \boldsymbol{\rho}) \\ \psi_3(\psi_1(\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}), \psi_2(\psi_1(\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho}), \mathbf{a}, \boldsymbol{\rho}), \boldsymbol{\rho}) \end{pmatrix}.$$

Inspecting the block updates (4.2), (4.3) and (4.6), we see that the submaps  $\psi_1$ ,  $\psi_2$  and  $\psi_3$  are each continuous. Therefore, the composition map  $\psi$  is also continuous. The second step is to verify that  $f(\mathbf{x}) > f(\psi(\mathbf{x}))$  if  $\mathbf{x} \neq \psi(\mathbf{x})$ . Consider the intermediate iterates

$$\begin{aligned}\boldsymbol{\beta}_0^+ &\equiv \psi_1((\boldsymbol{\beta}_0, \mathbf{a}, \boldsymbol{\rho})), \\ \mathbf{a}^+ &\equiv \psi_2((\boldsymbol{\beta}_0^+, \mathbf{a}, \boldsymbol{\rho})), \\ \boldsymbol{\rho}^+ &\equiv \psi_3((\boldsymbol{\beta}_0^+, \mathbf{a}^+, \boldsymbol{\rho})).\end{aligned}$$

For any  $\mathbf{x} \in S$ , we have

$$f(\mathbf{x}) \geq f((\boldsymbol{\beta}_0^+, \mathbf{a}, \boldsymbol{\rho})) \geq f((\boldsymbol{\beta}_0^+, \mathbf{a}^+, \boldsymbol{\rho})) \geq f(\mathbf{x}^+).$$

If, however,  $\mathbf{x}$  is not a fixed point of  $\psi$ , namely,  $\psi(\mathbf{x}) \neq \mathbf{x}$ , then at least one of the above inequalities is strict and, therefore,  $f(\mathbf{x}) > f(\psi(\mathbf{x}))$ . The third step is

to identify an  $\mathbf{x}_0 \in S$  so that the set  $\mathcal{L}_f(\mathbf{x}_0)$  is compact. Note that  $f$  blows up for  $\|\mathbf{a}\| \rightarrow \infty$  or  $\|\boldsymbol{\beta}_0\| \rightarrow \infty$ . Therefore, the set  $\mathcal{L}_f(\mathbf{x}_0)$  is compact for any initial  $\mathbf{x}_0 \in S$ . By Theorem C.1, it follows that the cluster points of the iterates of Algorithm 2 are fixed points of the algorithm map  $\psi(\mathbf{x})$  corresponding to Algorithm 2. To complete the proof, we just need to show that every fixed point of  $\psi(\mathbf{x})$  is a stationary point of the PNLL (4.1). If  $\mathbf{x}$  is a fixed point of  $\psi$ , then

$$\begin{aligned}\frac{\partial}{\partial \boldsymbol{\beta}_0} f(\mathbf{x}) &= \mathbf{0}, \\ \frac{\partial}{\partial \mathbf{a}} f(\mathbf{x}) &= \mathbf{0}, \\ \frac{\partial}{\partial \boldsymbol{\rho}} [g(\boldsymbol{\rho}|\boldsymbol{\rho}) + \lambda_2 J_2(\boldsymbol{\rho})] &= \mathbf{0}.\end{aligned}$$

In light of (B.4), it is clear that the last condition is equivalent to

$$\frac{\partial}{\partial \boldsymbol{\rho}} f(\mathbf{x}) = \mathbf{0}.$$

Therefore, every cluster point of the iterate sequence generated by Algorithm 2 is a stationary point of the PNLL (4.1).  $\square$

**Acknowledgments.** The authors thank John Treichler and Kathy Ensor for helpful discussions.

## REFERENCES

- ALLEN, G. I. (2013). Sparse and functional principal components analysis. Available at [arXiv:1309.2895](https://arxiv.org/abs/1309.2895).
- ALLEN, G. I., GROSENICK, L. and TAYLOR, J. (2014). A generalized least-square matrix decomposition. *J. Amer. Statist. Assoc.* **109** 145–159. [MR3180553](https://doi.org/10.1198/016214513012180553)
- BECKER, M. P., YANG, I. and LANGE, K. (1997). EM algorithms without missing data. *Stat. Methods Med. Res.* **6** 38–54.
- BÖHNING, D. and LINDSAY, B. G. (1988). Monotonicity of quadratic-approximation algorithms. *Ann. Inst. Statist. Math.* **40** 641–663. [MR0996690](https://doi.org/10.1080/03610928808839690)
- CHUNG, F. R. K. (1997). *Spectral Graph Theory*. *CBMS Regional Conference Series in Mathematics* **92**. Amer. Math. Soc., Providence, RI. [MR1421568](https://doi.org/10.1090/S0895-4713-1997-0883969-0)
- DEBOSSCHER, J., SARRO, L., LÓPEZ, M., DELEUIL, M., AERTS, C., AUVERGNE, M., BAGLIN, A., BAUDIN, F., CHADID, M., CHARPINET, S. et al. (2009). Automated supervised classification of variable stars in the CoRoT programme. *Astron. Astrophys.* **506** 519.
- DUBATH, P., RIMOLDINI, L., SÜVEGES, M., BLOMME, J., LÓPEZ, M., SARRO, L. M., DE RIDDER, J., CUYPERS, J., GUY, L., LECOEUR, I. et al. (2011). Random forest automated supervised classification of Hipparcos periodic variable stars. *Mon. Not. R. Astron. Soc.* **414** 2602–2617.
- FREEDMAN, W. L. (1988). New Cepheid distances to nearby galaxies based on BVRI CCD photometry. I. IC 1613. *Astrophys. J.* **326** 691–709.
- FRIEDMAN, J. H. (1984). A variable span smoother. Technical report, Stanford Univ., Stanford, CA.
- GELFAND, A. E. and VOUNATSOU, P. (2003). Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* **4** 11–15.

- GRAHAM, M. J., DRAKE, A. J., DJORGOVSKI, S. G., MAHABAL, A. A., DONALEK, C., DUAN, V. and MAKER, A. (2013). A comparison of period finding algorithms. *Mon. Not. R. Astron. Soc.* **434** 3423–3444.
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. Springer, New York. [MR2722294](#)
- HOERL, A. E. and KENNARD, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12** 55–67.
- HUANG, J. Z., SHEN, H. and BUJA, A. (2009). The analysis of two-way functional data using two-way regularized singular value decompositions. *J. Amer. Statist. Assoc.* **104** 1609–1620. [MR2750581](#)
- IVEZIĆ, Ž., SMITH, J. A., MIKNAITIS, G., LIN, H., TUCKER, D., LUPTON, R. H., GUNN, J. E., KNAPP, G. R., STRAUSS, M. A., SESAR, B. et al. (2007). Sloan digital sky survey standard star catalog for stripe 82: The dawn of industrial 1% optical photometry. *Astron. J.* **134** 973.
- KELLER, J. P., OLIVES, C., KIM, S. Y., SHEPPARD, L., SAMPSON, P. D., SZPIRO, A. A., ORON, A. P., LINDSTRÖM, J., VEDAL, S. and KAUFMAN, J. D. (2015). A unified spatiotemporal modeling approach for predicting concentrations of multiple air pollutants in the Multi-Ethnic Study of Atherosclerosis and Air Pollution. *Environ. Health Perspect.* **123** 301–309.
- LANGE, K., HUNTER, D. R. and YANG, I. (2000). Optimization transfer using surrogate objective functions. *J. Comput. Graph. Statist.* **9** 1–59. [MR1819865](#)
- LI, C. and LI, H. (2008). Network-constrained regularization and variable selection for analysis of genomic data. *Bioinformatics* **24** 1175–1182.
- LOMB, N. R. (1976). Least-squares frequency analysis of unequally spaced data. *Astrophys. Space Sci.* **39** 447–462.
- LONG, J. P., CHI, E. C. and BARANIUK, R. G. (2014). Estimating a common period for a set of irregularly sampled functions with applications to periodic variable star data. Preprint. Available at [arXiv:1412.6520](#).
- MEYER, R. R. (1976). Sufficient conditions for the convergence of monotonic mathematical programming algorithms. *J. Comput. System Sci.* **12** 108–121. [MR0439211](#)
- PERCY, J. R. (2007). *Understanding Variable Stars*. Cambridge Univ. Press, Cambridge.
- RAMSAY, J. O. (2004). Functional Data Analysis. In *Encyclopedia of Statistical Sciences*. Wiley, New York.
- RAÑOLA, J. M., NOVEMBRE, J. and LANGE, K. (2014). Fast spatial ancestry via flexible allele frequency surfaces. *Bioinformatics* **30** 2915–2922.
- REIMANN, J. D. (1994). Frequency estimation using unequally-spaced astronomical data. Ph.D. thesis, Univ. California.
- RICHARDS, J. W., STARR, D. L., BUTLER, N. R., BLOOM, J. S., BREWER, J. M., CRELLIN-QUICK, A., HIGGINS, J., KENNEDY, R. and RISCHARD, M. (2011). On machine-learned classification of variable stars with sparse and noisy time-series data. *Astrophys. J.* **733** 10.
- RIESS, A. G., MACRI, L., CASERTANO, S., LAMPEITL, H., FERGUSON, H. C., FILIPPENKO, A. V., JHA, S. W., LI, W. and CHORNOCK, R. (2011). A 3% solution: Determination of the Hubble constant with the Hubble space telescope and wide field camera 3. *Astrophys. J.* **730** 119.
- SCARGLE, J. D. (1982). Studies in astronomical time series analysis. II-Statistical aspects of spectral analysis of unevenly spaced data. *Astrophys. J.* **263** 835–853.
- SCHLAFLY, E. F., FINKBEINER, D. P., JURÍĆ, M., MAGNIER, E. A., BURGETT, W. S., CHAMBERS, K. C., GRAV, T., HODAPP, K. W., KAISER, N., KUDRITZKI, R.-P. et al. (2012). Photometric calibration of the first 1.5 years of the Pan-STARRS1 survey. *Astrophys. J.* **756** 158.
- SESAR, B., IVEZIĆ, Ž., LUPTON, R. H., JURÍĆ, M., GUNN, J. E., KNAPP, G. R., DE LEE, N., SMITH, J. A., MIKNAITIS, G., LIN, H. et al. (2007). Exploring the variable sky with the sloan digital sky survey. *Astron. J.* **134** 2236.



- SE SAR, B., IVEZIĆ, Ž., GRAMMER, S. H., MORGAN, D. P., BECKER, A. C., JURIC, M., DE LEE, N., ANNIS, J., BEERS, T. C., FAN, X. et al. (2010). Light curve templates and galactic distribution of RR Lyrae stars from Sloan Digital Sky Survey Stripe 82. *Astrophys. J.* **708** 717.
- SHAPPEE, B. J. and STANEK, K. Z. (2011). A new Cepheid distance to the giant spiral M101 based on image subtraction of Hubble Space Telescope/Advanced Camera for Surveys observations. *Astrophys. J.* **733** 124.
- STETSON, P. B. (1996). On the automatic determination of light-curve parameters for Cepheid variables. *Publ. Astron. Soc. Pac.* **108** 851–876.
- SUVEGES, M., SESAR, B., VARADI, M., MOWLAVI, N., BECKER, A. C., IVEZIĆ, Ž. BECK, M., NIENARTOWICZ, K., RIMOLDINI, L., DUBATH, P. et al. (2012). Search for high-amplitude Scuti and RR Lyrae stars in Sloan Digital Sky Survey Stripe 82 using principal component analysis. *Mon. Not. R. Astron. Soc.* **424** 2528–2550.
- TIAN, T. S., HUANG, J. Z., SHEN, H. and LI, Z. (2012). A two-way regularization method for MEG source reconstruction. *Ann. Appl. Stat.* **6** 1021–1046. [MR3012519](#)
- TORABI, M. and ROSYCHUK, R. J. (2010). Spatio-temporal modelling of disease mapping of rates. *Canad. J. Statist.* **38** 698–715. [MR2753010](#)
- UDALSKI, A., SZYMANSKI, M. K., SOSZYNSKI, I. and POLESKI, R. (2008). The Optical Gravitational Lensing Experiment. Final reductions of the OGLE-III data. *Acta Astron.* **58** 69–87.
- VANDERPLAS, J. T. and IVEZIĆ, Ž. (2015). Periodograms for multiband astronomical time series. Preprint. Available at [arXiv:1502.01344](#).
- WATKINS, L. L., EVANS, N. W., BELOKUROV, V., SMITH, M. C., HEWETT, P. C., BRAMICH, D. M., GILMORE, G. F., IRWIN, M. J., VIDRIH, S., ZUCKER, D. B. et al. (2009). Substructure revealed by RR Lyraes in SDSS Stripe 82. *Mon. Not. R. Astron. Soc.* **398** 1757–1770.
- WELCH, D. L. and STETSON, P. B. (1993). Robust variable star detection techniques suitable for automated searches—New results for NGC 1866. *Astron. J.* **105** 1813–1821.
- ZECHMEISTER, M. and KÜRSTER, M. (2009). The generalised Lomb–Scargle periodogram. A new formalism for the floating-mean and Keplerian periodograms. *Astron. Astrophys.* **496** 577–584.

J. P. LONG  
DEPARTMENT OF STATISTICS  
TEXAS A&M UNIVERSITY  
COLLEGE STATION, TEXAS 77843  
USA  
E-MAIL: [jljong@stat.tamu.edu](mailto:jljong@stat.tamu.edu)

E. C. CHI  
DEPARTMENT OF STATISTICS  
NORTH CAROLINA STATE UNIVERSITY  
RALEIGH, NORTH CAROLINA 27607  
USA  
E-MAIL: [eric\\_chi@ncsu.edu](mailto:eric_chi@ncsu.edu)

R. G. BARANIUK  
DEPARTMENT OF ELECTRICAL  
AND COMPUTER ENGINEERING  
RICE UNIVERSITY  
HOUSTON, TEXAS 77005  
USA  
E-MAIL: [richb@rice.edu](mailto:richb@rice.edu)