

## A Gaussian Mixture Model for Nulling Pulsars

D. L. KAPLAN,<sup>1</sup> J. K. SWIGGUM,<sup>1</sup> T. D. J. FICHTENBAUER,<sup>1</sup> AND M. VALLISNERI<sup>2,3</sup>

<sup>1</sup>*Center for Gravitation, Cosmology and Astrophysics, Department of Physics, University of Wisconsin–Milwaukee, P.O. Box 413, Milwaukee, WI 53201, USA*

<sup>2</sup>*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA*

<sup>3</sup>*TAPIR, California Institute of Technology, Pasadena, CA 91125, USA*

(Accepted January 27, 2018)

Submitted to ApJ

### ABSTRACT

The phenomenon of pulsar nulling – where pulsars occasionally turn off for one or more pulses – provides insight into pulsar-emission mechanisms and the processes by which pulsars turn off when they cross the “death line.” However, while ever more pulsars are found that exhibit nulling behavior, the statistical techniques used to measure nulling are biased, with limited utility and precision. In this paper we introduce an improved algorithm, based on Gaussian mixture models, for measuring pulsar nulling behavior. We demonstrate this algorithm on a number of pulsars observed as part of a larger sample of nulling pulsars, and show that it performs considerably better than existing techniques, yielding better precision and no bias. We further validate our algorithm on simulated data. Our algorithm is widely applicable to a large number of pulsars even if they do not show obvious nulls. Moreover, it can be used to derive nulling probabilities of nulling for individual pulses, which can be used for in-depth studies.

*Keywords:* methods — data analysis, methods — statistical, pulsars — general

### 1. INTRODUCTION

Pulsar surveys to date have found  $\sim 2,700$  pulsars (Manchester et al. 2016)<sup>1</sup>, most of which move across the  $P-\dot{P}$  diagram and turn off in  $\sim 10^7$  yr when they pass the “death line” (see, e.g., Lorimer & Kramer 2012). Just how and why pulsars turn off is still a subject of intense scrutiny, with ongoing observational and theoretical investigations (see references below). Nulling pulsars (Backer 1970) – pulsars whose radio emission ceases temporarily for one or more rotations – offer an invaluable laboratory to study pulsar emission mechanisms and magnetospheres. The meaning of null durations, the intervals between them, and the underlying mechanism have been matters of debate ever since the behavior was recognized.

The first comprehensive study of nulling pulsars (Ritchings 1976) suggested that as a pulsar ages, the

time interval between regular bursts of pulsed emission increases, eventually leading to “death” when the interval between bursts is much longer than the duration of the bursts themselves. A later study of 72 pulsars found a stronger correlation between null fraction (NF; the fraction of time that a pulsar spends in a null state) and spin period, but still argued that nulling could be indicative of a faltering emission mechanism (Biggs 1992). Wang et al. (2007) argued on the basis of a smaller sample of 45 nulling pulsars that nulling behavior was more related to a large characteristic age than other parameters, but the analysis was not quantitative.

However, while data-collection capabilities, processing techniques, and theories explaining pulsar nulling (e.g., Gajjar et al. 2012) have become more sophisticated, we are still in a regime where increasing the sample size of nullers along with the precision of the nulling analysis can have a significant impact on our understanding. We are working to increase the sample size through detailed followup observations of sources found in the Green Bank North Celestial Cap survey (GBNCC; Stovall et al. 2014; Lynch et al. 2018; Kawash et al. 2018).

Corresponding author: D. L. Kaplan  
kaplan@uwm.edu

<sup>1</sup> <http://www.atnf.csiro.au/people/pulsar/psrcat>

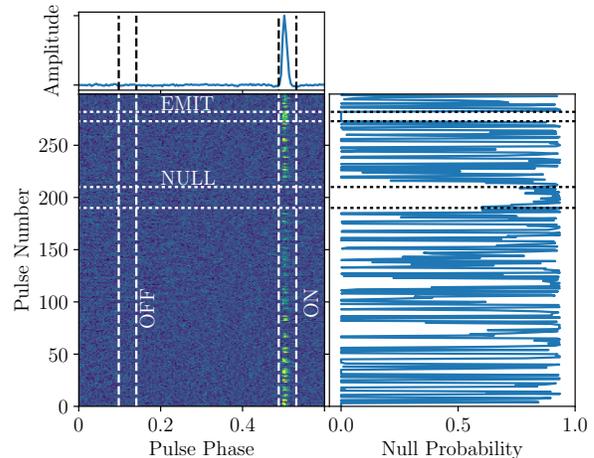
At the same time, the classification of pulsars into those that do or do not null is largely based on qualitative examination by eye, and computation of the nulling fraction suffers from significant limitations that may bias the results.

Here we present a robust technique for determining nulling fractions using Gaussian mixture models (see, e.g., Ivezic et al. 2014). As we demonstrate, the robustness of this technique allows us to infer quantitative limits on nulling fractions even for pulsars that have no obvious nulls, so the technique can be used for more sophisticated population analyses. First, we describe the sample data that we used to test our technique (§ 2). We then discuss the technique itself (§ 3), along with fit results to actual data (§ 3.2) and simulated data (§ 3.3). We conclude in § 4. All of our source code is available at [10.5281/zenodo.1155855](https://zenodo.org/record/1155855). Note that we believe a similar algorithm was applied to nulling pulsars by Arjunwadkar et al. (2014), but we have been unable to find details of its implementation or results.

## 2. SAMPLE DATA

To test our method, we gathered timing data from a recent study of new pulsar discoveries in the GBNCC survey (Lynch et al. 2018). Data were collected with the 100-m Robert C. Byrd Green Bank Telescope (GBT), observing at a center frequency of 800 MHz. We digitized a 200 MHz bandwidth using the Green Bank Ultimate Pulsar Processing Instrument (GUPPI; DuPlain et al. 2008) in incoherent search mode, with 2,048 frequency channels and sampling every 40.96  $\mu$ s. Pulsars included in this study (PSRs J0323+6742, J0054+6946, and J0137+6349) were observed monthly for 3-minute exposures between February 2013 and January 2014. Exposure times for PSR J0054+6946 were lengthened to 6 minutes starting in November 2013 through May 2014.

For each pulsar, data were folded modulo the pulse period using DSPSR and RFI was zapped interactively using the `pazi` routine from PSRCHIVE. A crucial step in nulling studies is the choice of *on-pulse* and *off-pulse* windows, defined as fixed phase intervals (within folded profiles) that are assumed to contain most and none, respectively, of the pulsar emission. Loosely speaking, the intensities observed in the off-pulse window are attributed to background noise setting the threshold for on-pulse-window intensities to be classified as nulling or emitting. We set the width and phase of on-pulse windows manually by inspecting folded profiles, choosing off-pulse windows of the same size at pulse phases containing no visible emission, often about 0.5 rotations



**Figure 1.** Amplitude versus pulse phase and pulse number for a subset of the data on PSR J0323+6742. The top panel shows the average pulse profile, with the on- and off-pulse phase windows delineated by the vertical lines. The bottom-left panel shows the amplitude for individual pulses, with the phase windows indicated. The right panel shows the probability of nulling (Eqn. 5). We highlight regions with significant nulling behavior (20 successive pulses with minimum *a posteriori* null probability greater than 62%) and relatively steady emission (10 pulses with maximum *a posteriori* null probability  $\sim$  0%) with the horizontal lines. Such periods of successive nulls or emitted pulses are not intrinsically different from those of pulsars that switch states more rapidly, but are easier to visualize.

from the on-pulse window, but careful not to include any interpulse (if present).

We fitted and removed a 6th-order polynomial to flatten the baseline of each single pulse profile and center the off-pulse noise on zero intensity (similar to Lynch et al. 2013; Rosen et al. 2013). Afterwards, we assembled our dataset of on and off intensities for each individual pulse by integrating over each phase bin and concatenating the results from the separate 3-min/6-min exposures. See Figure 1 for an illustration of the pulse intensities measured in on- and off-pulse windows, along with some significant nulling behavior.

## 3. A GAUSSIAN MIXTURE MODEL FOR NULLING PULSARS

The standard algorithm for fitting the nulling fraction NF (Ritchings 1976) is to first construct histograms of the integrated intensities  $I^{\text{ON}} \equiv \{I_k^{\text{ON}}\}$  and  $I^{\text{OFF}} \equiv \{I_k^{\text{OFF}}\}$  measured in the on-pulse and off-pulse windows at each observation  $k$ . We denote the histograms as  $\text{ON}_n(I^{\text{ON}})$  and  $\text{OFF}_n(I^{\text{OFF}})$ , where the index  $n$  identifies histogram bins. Then, for a series of trial values of NF one computes the differ-

ence histogram  $\Delta_n(I^{\text{ON}}, I^{\text{OFF}}) = \text{ON}_n(I^{\text{ON}}) - \text{NF} \times \text{OFF}_n(I^{\text{OFF}})$ . The best-fit  $\hat{\text{NF}}$  is the value that minimizes the sum of  $\Delta_n(I^{\text{ON}}, I^{\text{OFF}})$  over negative intensities,  $|\sum_{I_n < 0} \Delta_n(I^{\text{ON}}, I^{\text{OFF}})|$ , where the nulling is presumed to dominate. This method has a number of drawbacks: first, it requires construction of histograms so there is an arbitrary choice of binning. Second, it assumes that the pulses with  $I_k^{\text{ON}} < 0$  are entirely due to nulling, excluding weak pulsars where emitted pulses be overwhelmed by radiometer noise or otherwise end up at negative intensities.

### 3.1. Fitting Algorithm

In our proposed method, we define the likelihood of the on-pulse dataset  $I^{\text{ON}}$  as a one-dimensional Gaussian mixture model,

$$p(I^{\text{ON}}|\{w_j, \mu_j, \sigma_j\}) = \prod_k^N \left( \sum_{j=1}^M w_j \mathcal{N}(I_k^{\text{ON}}|\mu_j, \sigma_j) \right), \quad (1)$$

where  $\mu_j$  and  $\sigma_j$  are the means and standard deviations of  $M$  normal distributions, and  $w_j$  are their weights in the mixture, which must satisfy  $\sum w_j = 1$ . Therefore there are  $3M - 1$  free parameters in this model. The case with  $M = 2$  corresponds to the standard description of nulling with two modes; we order results so that the  $j = 1$  component has the lowest mean and describes the nulls, so  $\text{NF} = w_1$ . Note that Eq. (1) implies that the error in the measurement of the intensities is negligible compared to the intrinsic scatter in the mixture components, described by the  $\sigma_j$ .<sup>2</sup>

The parameters that maximize this likelihood can be found using the expectation-maximization algorithm (Dempster et al. 1977), implemented as `GaussianMixture` in `scikit-learn` (Pedregosa et al. 2011). Using this implementation, we find reasonable results for pulsars with significant separations between nulling and emitting pulses. Such separation may be a signal-to-noise effect, but scatter in pulse intensities could also result from intrinsic variability in the pulsars themselves or interstellar scintillation (Rickett 1990; Jenet & Gil 2003). However, as we show below the results are not ideal for other pulsars.

We can do better by using the off-pulse dataset  $I^{\text{OFF}}$  to constrain the null-component parameters  $\mu_1$  and  $\sigma_1$ , by way of the off-pulse likelihood

$$p(I^{\text{OFF}}|\mu_1, \sigma_1) = \prod_k \mathcal{N}(I_k^{\text{OFF}}|\mu_1, \sigma_1). \quad (2)$$

<sup>2</sup> If the measurement error is not negligible, the  $\sigma_j$  are effectively redefined to include it, under the assumptions that the error is similar for every observation.

Indeed, we may think of this step as providing a prior distribution for  $\mu_1$  and  $\sigma_1$ , which is then used in Eqn. 1. We then explore the  $\{w_j, \mu_j, \sigma_j\}$  parameter space using Markov Chain Monte Carlo (MCMC) techniques, specifically the affine-invariant population-MCMC algorithm `emcee` described by Foreman-Mackey et al. (2013). The details of our implementation are as follows:

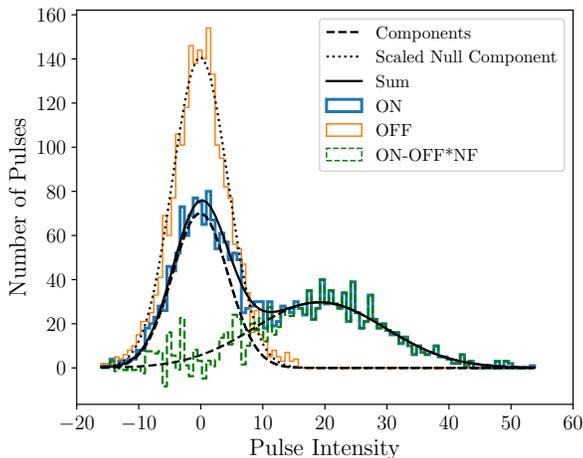
- We initialize 40 `emcee` “walkers” around the best-fit region for  $w_j, \mu_j, \sigma_j$ , as determined by expectation maximization run on  $I_k^{\text{ON}}$ .
- For simplicity, we set the  $\mu_1$  and  $\sigma_1$  priors as normal distributions centered on  $\frac{1}{k} \sum_k I_k^{\text{OFF}}$  and on the inner-quartile range of the  $I^{\text{OFF}}$ , respectively, with widths determined following Ahn & Fessler (2003).
- Prior distributions for  $\mu_j$  and  $\sigma_j$  (with  $j > 1$ ) are taken to be flat. The prior distribution for the weights  $w_j$  is a Dirichlet distribution, but since the sum of the weights is 1 it is effectively flat.
- We run the walkers through 50 iterations to achieve “burn in.”
- Last, we run the walkers for 500 iterations to obtain the final population, representative of the  $\mu_j, \sigma_j$ , and  $w_j$  joint posterior.

We experimented with increasing the number of walkers and iterations and found that the values above gave sufficiently reliable results for data-sets with a few thousand pulses and  $M < 5$ , but they can be increased as needed to achieve reliable posterior distributions.

### 3.2. Fit Results

Representative results from this algorithm are shown in Figure 2, where we plot  $p(I^{\text{ON}}|\{\hat{w}_j, \hat{\mu}_j, \hat{\sigma}_j\})$  (solid line) and  $p(I^{\text{OFF}}|\hat{\mu}_1, \hat{\sigma}_1)$  (dashed) on top of the  $I^{\text{ON}}$  and  $I^{\text{OFF}}$  histograms (with appropriate normalizations); here  $\{\hat{w}_j, \hat{\mu}_j, \hat{\sigma}_j\}$  are the *a posteriori* joint maxima of the Gaussian mixture parameters. Figure 3 shows the posterior densities of the Gaussian means and variances. The data appear to be fit well, with the null component ending up close to the off-pulse fit results; we observe no significant pathologies in the MCMC chains. To evaluate goodness of fit quantitatively, we perform the Kolmogorov–Smirnov test (KS test; Chakravarti et al. 1967), and find a statistic value of 0.015, corresponding to a  $p$ -value of 0.7—no evidence to reject the hypothesis that the data were sampled from the best-fit distribution.

For this pulsar we find that the emitting and the null components are sufficiently separate that all of the algorithms outlined above would give similar results. For

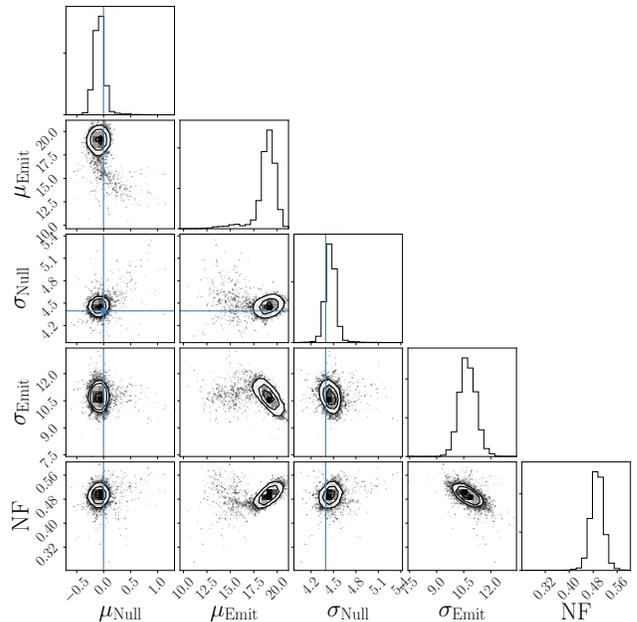


**Figure 2.** Distribution of pulse intensities for PSR J0323 + 6742. The blue and orange histograms are the raw intensities for on- and off-pulse windows, respectively. The dashed curves are the maximum *a posteriori* individual components from the Gaussian mixture model for  $M = 2$ , with the solid curve their sum, as determined by our MCMC algorithm. The dotted curve is the component for the nulls scaled by  $1/\text{NF}$ : it matches the off-pulse intensities well. Finally, the dashed green histogram is the data that would be used to implement the [Ritchings \(1976\)](#) algorithm, although it is plotted for our best-fit value of  $\text{NF}$ . Using our MCMC algorithm we find  $\text{NF} = 50 \pm 2\%$ , compared to 56% using [Ritchings \(1976\)](#).

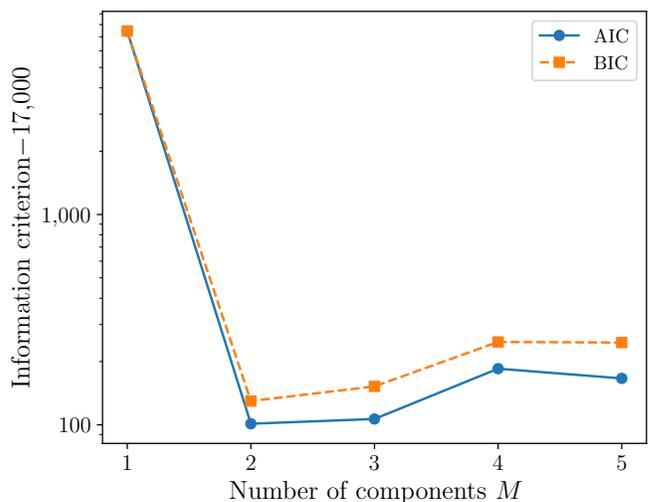
instance, we find that only about 2% of the pulses from the emitting component would have intensities less than 0, which would only bias the [Ritchings \(1976\)](#) results by a small amount.

If the problem is well behaved, we can select the optimal number of components  $M$  by maximizing the Bayesian information criterion (BIC; [Schwarz 1978](#)) or the Akaike information criterion (AIC; [Akaike 1974](#)). We illustrate this in [Figure 4](#), which shows both. We find strong evidence that nulling behavior is present ( $M > 1$ ), and a weaker preference for  $M = 2$  compared to  $M > 2$ . The BIC corresponds to approximating the Bayes ratios between models as  $e^{-\Delta_{\text{BIC}}/2}$ , where  $\Delta_{\text{BIC}}$  is the difference in BIC (likewise for AIC). In this approximation, the implied Bayes ratio for  $M = 2$  vs.  $M = 3$  is  $O_{23} = 175$  (AIC) or  $O_{23} = 92,000$  (BIC), showing that  $M = 2$  is indeed preferred.

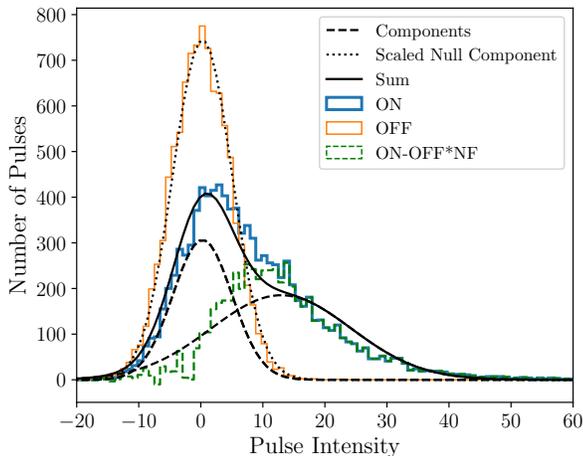
We show another example in [Figure 5](#), where for  $M = 2$  we find on-pulse maximum *a posteriori* parameters  $\hat{\mu}_2 = 12.9$  and  $\hat{\sigma}_2 = 11.3$ , with  $\hat{\text{NF}} = 41\%$ . In this case there is much less separation in intensities between nulls and pulses, with about 13% of the pulses having  $I < 0$ . Nonetheless our method gives a robust fit. [Figure 5](#) appears to show slight deviations from Gaussian



**Figure 3.** Posterior probability densities for the Gaussian-distribution parameters of the null and emitting components, as derived in our MCMC algorithm. The solid lines show the  $\mu_1$  and  $\sigma_1$  modes as inferred from the off-pulse data. The contours are 1-, 2-, and 3- $\sigma$  joint confidence contours.



**Figure 4.** Model-comparison evidence for different number of Gaussian-mixture components in the PSR J0323 + 6742 data. We show the Bayesian information criterion (BIC; orange squares) and the Akaike information criterion (AIC; blue circles) as a function of the number of components  $M$ . The non-nulling hypothesis ( $M = 1$ ) is rejected at high confidence, and  $M = 2$  is preferred by both criteria.



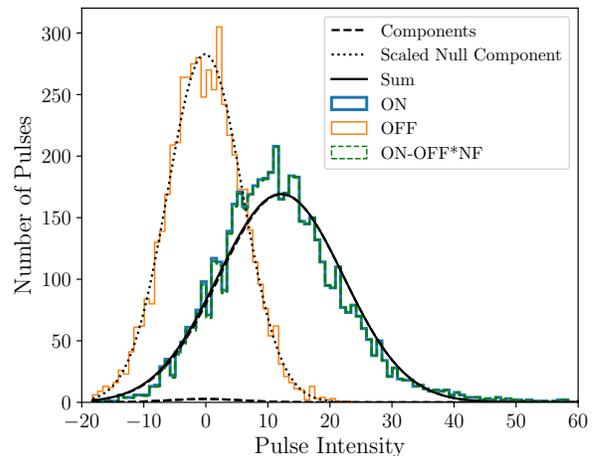
**Figure 5.** Distribution of pulse intensities for PSR J0054 + 6946 (see Fig. 2 for details). Here the nulling and emitting components are much less separated than in Figure 3. Using our MCMC algorithm we find  $\hat{NF} = 40.8 \pm 1.4\%$ , compared to 48% using [Ritchings \(1976\)](#).

distributions, which may be handled with more complex models that (for example) incorporate asymmetric distributions due to scintillation. Indeed, the KS test rejects the assumption that the data were drawn from the best-fit distribution at the  $5 \times 10^{-12}$  level; nevertheless, since our goal is primarily to quantify the bimodality of the emitted pulses rather than the exact intensity distribution, we believe the nulling results themselves to be robust.

Finally, in Figure 6 we show an example where the pulsar shows no obvious nulling. Our algorithm finds  $\hat{NF} = 0.9 \pm 6.4\%$ , consistent with 0, but the [Ritchings \(1976\)](#) algorithm still returns a non-zero value of 21%. Again we see deviations from Gaussian distributions (with a KS-test  $P$  value of  $2 \times 10^{-7}$ ), but the overall robustness of our determination is evident. In contrast to [Ritchings \(1976\)](#) which can give determinations of non-zero nulling fractions even for pulsars that do not appear to null, our algorithm behaves well. Therefore we can use it for all pulsars that are sufficiently bright regardless of whether nulling is evident, and derive more robust determinations of whether or not weak nulling behavior is present.

### 3.3. Simulation Results

We validate our algorithm by simulating pulsar data for a range of parameters, drawing random intensities according to Eqs. (2) and (1) for the off- and on-pulse windows, respectively. We base our synthetic datasets on the PSR J0323 + 6742 data analyzed above: we simulate 2,000 pulses with  $\mu_1 = 0$ ,  $\sigma_1 = 5$ ,  $\sigma_2 = 10$ , and



**Figure 6.** Distribution of pulse intensities for PSR J0137 + 6349 (see Fig. 2 for details). Here there is no obvious evidence for nulling. Using our MCMC algorithm we find  $\hat{NF} = 0.9 \pm 6.4\%$ , compared to 21% using [Ritchings \(1976\)](#).

nulling fraction  $NF = 0.5$ . We vary  $\mu_2$  between 5 (hard to distinguish from the nulls) and 30 (easily distinguishable), and we repeat the test for 30 trials for each value of  $\mu_1$ .

We plot the median and standard deviation of the estimated NF in Figure 7, using the [Ritchings \(1976\)](#) method, the expectation-maximization method, and our Bayesian algorithm (in which case we report the maximum *a posteriori* NF). All three algorithms agree for high pulse intensities,  $\mu_2 \gtrsim 25$ .

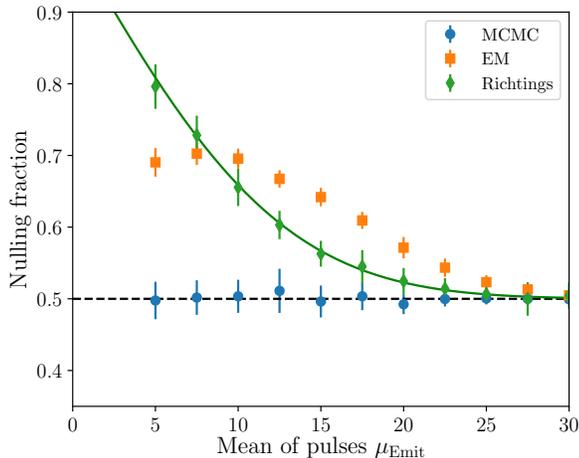
We see that the [Ritchings \(1976\)](#) NF estimates are highly biased for low values of  $\mu_2$ , as expected. This is because the model has a significant fraction of non-nulled pulses with intensities less than 0, ranging from 30% (for  $\mu_2 = 5$ ) to 0.1% (for  $\mu_2 = 30$ ). Specifically, we expect a fraction

$$\frac{1}{2} - \frac{1}{2} \operatorname{erf} \left( \frac{\mu_2 \sqrt{2}}{2\sigma_2} \right) \quad (3)$$

of the emitted pulses to have intensities  $< 0$ , where  $\operatorname{erf}(x)$  is the error function of  $x$ . This then leads to a biased estimated NF,

$$NF + (1 - NF) \left[ 1 - \operatorname{erf} \left( \frac{\mu_2 \sqrt{2}}{2\sigma_2} \right) \right] \quad (4)$$

which we have plotted in Figure 7, where they agree with our simulated results. The EM results using `GaussianMixture` are also biased at low pulse intensities. By contrast, our Bayesian algorithm performs well, with consistent uncertainties and no obvious bias across the  $\mu_2$  range.



**Figure 7.** Comparison on NF estimates for simulated data, as derived using the Ritchings (1976) algorithm (green diamonds), our expectation-maximization algorithm (orange squares, labeled as “EM”), and our Bayesian algorithm fit (blue circles). The horizontal line marks the true nulling fraction of 0.5. Each simulated dataset consisted of 2,000 pulses, with  $\mu_1 = 0$ ,  $\sigma_1 = 5$ ,  $\sigma_2 = 10$ ; simulations were repeated 30 times for each value of  $\mu_2$  (hence the vertical error bars). The solid green curve shows our analytical expectation for the bias of the Ritchings (1976) algorithm (see main text).

#### 4. DISCUSSION AND CONCLUSIONS

We have outlined and demonstrated an improved method to determine the nulling fraction of a pulsar. The method performs well in the limit of weak nulling, so it can be applied to a large number of pulsars without evident strong nulls. Unlike the traditional Ritchings (1976) algorithm, our method is unbiased, and it can be applied to pulsars with more than two emission modes, as long as those are reflected in the pulse intensities. However, it does require specification of the functional form of the intensity distributions for the nulling and emitting components: here we assume sums of Gaussians, although exponentials appropriate for 100% modulation by interstellar scintillation (e.g., Rickett 1990), or intermediate distributions are also possible. In those cases the AIC/BIC values can be used to quantitatively compare how well alternative distributions fit the data.

An additional benefit to this analysis is that we can determine explicitly the probability that any individual pulse belongs to a given class. This is sometimes called the “responsibility” (Hastie et al. 2009), and is given by:

$$p(j|I_k^{\text{ON}}) = \frac{w_j \mathcal{N}(\mu_j, \sigma_j)}{\sum_{i=1}^M w_i \mathcal{N}(\mu_i, \sigma_i)} \quad (5)$$

for class  $j$ . An example of this is shown in Figure 1, where we can determine the nulling probability as  $p(j = 1|I_k^{\text{ON}})$ . This probability can be computed for the maximum *a posteriori*  $\{\hat{w}_j, \hat{\mu}_j, \hat{\sigma}_j\}$ , or it can be marginalized over their distributions. Individual-pulse nulling probabilities can be used in robust multi-wavelength studies, to establish whether the X-ray properties of the pulses received during nulls differ from the others (e.g., Hermsen et al. 2013). We can also look for temporal patterns in the nulling properties, like the length of nulls (Fig. 1) or the time between nulls (e.g., Wang et al. 2007) using quantitative probability thresholds, and we could examine the probability that adjacent pulses transition between nulling and emitting behavior. Finally, we can fit for more than two components and identify mode changing quantitatively, in addition to nulling. All of these topics will be explored in future papers.

*Acknowledgments.* We thank S. McSweeney for helpful comments. We thank an anonymous referee and the AAS journals statistics editor for their suggestions. The Green Bank Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc. Support was provided by the NANOGrav NSF Physics Frontiers Center award number 1430284. MV acknowledges support from the JPL RTD program. Portions of this research were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

*Facility:* GBT

*Software:* Astropy (Astropy Collaboration et al. 2013), DSPSR (<http://dspsr.sourceforge.net/index.shtml>), emcee (Foreman-Mackey et al. 2013), PSRCHIVE (Hotan et al. 2004), scikit-learn (Pedregosa et al. 2011).

#### REFERENCES

- Ahn, S., & Fessler, J. A. 2003, Standard errors of mean, variance, and standard deviation estimators, Tech. Rep. Technical Report 413, Communications and Signal Processing Lab., Univ. of Michigan
- Akaike, H. 1974, IEEE Transactions on Automatic Control, 19, 716
- Arjunwadkar, M., Rajwade, K., & Gupta, Y. 2014, in Astronomical Society of India Conference Series, Vol. 13, 79–81

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
- Backer, D. C. 1970, *Nature*, 228, 42
- Biggs, J. D. 1992, *ApJ*, 394, 574
- Chakravarti, I. M., Laha, R. G., & Roy, J. 1967, *Handbook of Methods of Applied Statistics*, Vol. 1 (Wiley and Sons), 392–394
- Dempster, A. P., Laird, N. M., & Rubin, D. B. 1977, *Journal of the Royal Statistical Society, Series B*, 39, 1
- DuPlain, R., Ransom, S., Demorest, P., et al. 2008, in *Proc. SPIE*, Vol. 7019, *Advanced Software and Control for Astronomy II*, 70191D
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
- Gajjar, V., Joshi, B. C., & Kramer, M. 2012, *MNRAS*, 424, 1197
- Hastie, T., Tibshirani, R., & Friedman, J. 2009, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second Edition, Springer Series in Statistics (Springer New York)
- Hermesen, W., Hessels, J. W. T., Kuiper, L., et al. 2013, *Science*, 339, 436
- Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, *PASA*, 21, 302
- Ivezić, Ž., Connelly, A. J., VanderPlas, J. T., & Gray, A. 2014, *Statistics, Data Mining, and Machine Learning in Astronomy*
- Jenet, F. A., & Gil, J. 2003, *ApJL*, 596, L215
- Kawash, A., et al. 2018, *ApJ*, submitted
- Lorimer, D. R., & Kramer, M. 2012, *Handbook of Pulsar Astronomy* (Cambridge, UK: Cambridge University Press)
- Lynch, R. S., Swiggum, J. K., Kondratiev, V. I., Kaplan, D. L., et al. 2018, *ApJ*, submitted
- Lynch, R. S., Boyles, J., Ransom, S. M., et al. 2013, *ApJ*, 763, 81
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2016, *VizieR Online Data Catalog*, 1
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *Journal of Machine Learning Research*, 12, 2825
- Rickett, B. J. 1990, *ARA&A*, 28, 561
- Ritchings, R. T. 1976, *MNRAS*, 176, 249
- Rosen, R., Swiggum, J., McLaughlin, M. A., et al. 2013, *ApJ*, 768, 85
- Schwarz, G. 1978, *The Annals of Statistics*, 6, 461
- Stovall, K., Lynch, R. S., Ransom, S. M., et al. 2014, *ApJ*, 791, 67
- Wang, N., Manchester, R. N., & Johnston, S. 2007, *MNRAS*, 377, 1383