Memorandum 002

Characterizing DM Events in the Interstellar Medium

Vivek Gupta
Jacob E. Turner
T. Joseph W. Lazio
Daniel R. Stinebring

2016 December 31

http://nanograv.org/

Contents

Abstract

The NANOGrav 9 year data release contains unusual deviations in dispersion measure variations caused by abrupt changes in the electron density of the interstellar medium along the line of sight. We employ a couple of techniques in Bayesian analysis and Markov Chain Monte Carlo sampling to develop statistical models for dispersion measure in current NANOGrav observations, as well as predict future trends. These models are designed to be integrated into a "Quick Look" program, the output of which can then be used to identify DM "events" in recently collected data. We also analyze the negative DM event which occurred along the line of sight to PSR J1713+0747, which we have determined to be the result of an electron-devoid region in the ISM approximately 1.64 AU in width transverse to the LOS and 10^4 AU along the LOS. We then propose a few simple geometrical structures which could approximate the shape of the electron-deficient region in the ISM. Finally, we discuss a few potential physical sources which might explain the existence of such a region.

1 Introduction

NANOGrav is a North American-based collaboration which carries out radio frequency pulsar timing observations with the hope of detecting low frequency gravitational waves from a set of millisecond pulsars [4]. These observations are carried out using the William E. Gordon Radio Telescope in Arecibo, Puerto Rico, and the Robert C. Byrd Green Bank Telescope in Green Bank West Virginia. Pulse arrival times are calculated up to nanosecond precision and searched for correlated variations originating from gravitational waves passing between Earth and the pulsars. Other variations in the arrival times of pulses can be observed due to a large number of factors including the relative motion of a pulsar with respect to Earth, frequency dependent delays due to inhomogeneous interplanetary and interstellar ionized medium along the line of sight, and intrinsic pulsar spin evolution. We hope to understand and correct for all non-gravitational wave sources that manifest themselves as variations in the TOAs. A crucial component of the success of pulsar timing arrays relies on the understanding of how the interstellar medium affects timing accuracy.

Radio beams have to travel long distances through the interstellar medium (henceforth referred to as the ISM) before they are detected by our radio telescopes. The ISM causes variations in the pulsar TOAs through many of mechanisms, with one of the most prominent being free electron dispersion. The free electrons and ions present in the ISM cause a frequency dependent increase in the travel time of the pulses, with lower frequencies experiencing larger time delays than higher frequencies. We quantify this dispersive delay through a quantity known as dispersion measure, which is defined as

$$\mathrm{DM} = \int_0^d n_e dl. \tag{1}$$

Dispersion measure represents the integrated column density of free electrons present along the line of sight to the pulsar, and has been found to vary systematically for most of the pulsars observed by NANOGrav. Additionally, we occasionally see abrupt changes in the DM, known as DM events, for which we are currently lacking physical explanations.

Although NANOGrav observes at relatively frequent intervals of 1 week–3 weeks, the collaboration chooses to publish a few years worth of data at once, thus delaying the availability of processed data. However, we would require an almost real-time alert system for potential DM events in order to schedule observations that would yield useful information. Such a system would require accurate measurements and models of DM and consistent updating with new measurements.

In §2, we describe the models used and methods employed to obtain best fitting models for the present data. Section 3 includes the results and a discussion thereof. In §4, we look at the physical characteristics of a previously observed DM event in PSR J1713+0747, as well as define constraints on the dimensions and shape of a hole in the ISM which would be required to explain the drop in the observed DM. In §5, we discuss the future implications of our work. Finally, in §6 we provide a summary of our work.

Throughout, in order to characterize the DM variations, we adopt the standard NANOGrav technique of using DMX, which is a piecewise linear fit of the variations in DM that results from fitting within the `TEMPO` and `TEMPO2` software packages. Additional information on DMX, its definition, and analysis is contained within the NANOGrav Nine-Year Data Release [4] and Lam et al. [9].

2 Models & Methods

Proper modeling of DM variations must account for minute changes in the ISM along the LOS, which, on large scales, is thought to exhibit Kolmogorov turbulence [12]. Our models were trained on the NANOGrav 9 year data set and tested for effectiveness via trend prediction of the NANOGrav 11 year data set. We based our models on the idea that changes in the DM will be dominated by the changes in the ISM between Earth and the pulsar. In particular, we consider the effects of the solar wind and a few correlated changes in the ISM across the LOS [9]. To ensure the consistency of our models during the introduction of future data, we obtained our model parameters through a Bayesian approach to model fitting. In short, this method of inference relies on new data to update prior knowledge for a given phenomena in order to more accurately predict the outcome of future occurences.

From a mathematical standpoint, let us consider a set of model parameters θ that are collectively represented by a prior distribution $\Pr(\theta)$. When new data D is introduced, we can update our current model by multiplying it by a likelihood distribution $\Pr(\mathrm{D}|\theta)$. By applying the product rule, we get

$$\Pr(\mathrm{D}|\theta)\Pr(\theta) = \Pr(\theta|\mathrm{D})\Pr(\mathrm{D}), \tag{2}$$

where $\Pr(\theta|\mathrm{D})$ is the resulting posterior distribution and $\Pr(\mathrm{D})$ is the evidence for the model. Dividing by the evidence, we get

$$\Pr(\theta|\mathrm{D}) = \frac{\Pr(\mathrm{D}|\theta)\Pr(\theta)}{\Pr(\mathrm{D})}, \tag{3}$$

more commonly known as Bayes' Theorem [7, Chapter 1.3].

In our approach to modeling DM variations, we used two methods based on Bayesian inference.

2.1 Simple Bayesian Analysis

The DMX variations follow visible trends which can be easily modeled using a linear term and a sinusoidal term having period close to one year. The linear term accounts for the motion of the pulsar towards or away from earth, while the sinusoidal term is present because of the changing effects of the solar wind as the earth revolves around the sun. Some pulsars require an additional sine term with a period much longer than one year in order to model larger scale variations. As a result, our models are based upon simple equations having one linear term and one or two sine terms with different periods. They are described by

$$\begin{aligned}\mathrm{DMX}(t) = b + mt + A_1 \sin\left(\frac{2\pi t}{P_1} + \phi_1\right) \\ + A_2 \sin\left(\frac{2\pi t}{P_2} + \phi_2\right).\end{aligned} \tag{4}$$

Fitting our models using Bayesian analysis involved defining prior, likelihood and posterior probablity functions, as well as setting up a sampler to call these functions recursively. We chose flat uniform priors, defining a certain tentative range for each parameter and using 1 as our prior probability if a parameter was within the given range, and 0, otherwise. In order to maximize the likelihood estimate for our data, we chose our likelihood function to be a Gaussian based on residuals R between the current model and the data:

$$\ln(\mathrm{Likelihood}) = -\Sigma \frac{R^2}{2\sigma^2}, \tag{5}$$

where σ is the uncertainty on each data point.

The posterior probability is then defined as the product of the prior and the likelihood:

$$\ln(\mathrm{Prob}) = \ln(\mathrm{Prior}) + \ln(\mathrm{Likelihood}). \tag{6}$$

We then sampled from a n-dimensional parameter space by implementing the Metropolis Hastings algorithm, which iteratively generated a sequence of random samples from the prior such that the distribution of the next sample was dependent only on the current value (thus turning the sequence of samples into a Markov chain) [10] [6]. If the prior distribution was well sampled, the posterior represented a distribution of the best fit models based on our likelihood function.

The mean of the resulting distribution was taken as the maximum likelihood estimate, i.e. the best fit model for our data, while the standard deviation of the distribution was considered to be the error in estimating the best fit model.

Once we obtained the models, our task of calculating the significance of variation of a new observation was relatively simple. Each data point was converted into a gaussian with mean equal to the data value and standard deviation equal to the errorbar of that point. Thus a distribution of data points was converted to a collection of gaussians representing the data. Subtracting the model distribution from this collection of gaussians gave us a distribution of the residuals, where the mean was the actual residual at that point, and the standard deviation was the uncertainty or the error in calculating that residual.

The significance of the deviation ($\sigma-$value) of a new observation from the general trend was calculated as the ratio of the difference between residual of the new observation and the overall mean of previous residuals, over the standard deviation of the overall residuals. This is equivalent to determining the number of standard deviations away a new data point is from the mean of previous residuals, and can be given by

$$\sigma-\text{value} = \frac{\text{New Residual-Mean(previous residuals)}}{\text{Std.Dev.(previous residuals)}}. \tag{7}$$

[1]https://github.com/jellis18/PTMCMCSampler

2.2 Gaussian Process Regression

In our second method, we modeled the DM variations with a linear and sine term to account for the pulsar's position relative to Earth and the solar wind, respectively. We also computed Bayes factors to determine whether the data preferred the addition of a quadratic term to account for stochastic processes [3] [8]. Thus, when considering a linear model M_1 and a quadratic model M_2, the dispersion measure taken at some day t can be described by

$$\mathrm{DM}(t) = nct^2 + mt + A\sin\left(\frac{2\pi t}{P} + \phi\right) + b,$$
$$n = \begin{cases} 1, & \text{if } 2\ln\left(\frac{\Pr(\mathrm{D}|\mathrm{M}_2)}{\Pr(\mathrm{D}|\mathrm{M}_1)}\right) \geq 6 \\ 0, & \text{if } 2\ln\left(\frac{\Pr(\mathrm{D}|\mathrm{M}_2)}{\Pr(\mathrm{D}|\mathrm{M}_1)}\right) < 6 \end{cases}. \tag{8}$$

In order to obtain the most accurate models of future variations, we then introduced a Gaussian random variable for each point in our data. The covariance between any two points t_i and t_j could then be represented by a matrix K_{ij} and equivalently described by a function $k(t_i, t_j)$ [14]. For our purposes, our covariance function was described by the Matèrn-$\frac{3}{2}$ Kernel as provided by the Gaussian process package George [1]

$$k(t_i, t_j) = \left(1 + \sqrt{3(t_i - t_j)^2}\right)e^{-\sqrt{3(t_i - t_j)^2}}, \tag{9}$$

as it closely resembles the -8/3 spectral index typically seen in DM variations.

The resulting log likelihood function is then

$$\begin{aligned} &\ln\Pr(\mathrm{D} - \mathrm{DM}(t)|\mathrm{GP}) = \\ &-\frac{1}{2}(\mathrm{D} - \mathrm{DM}(t))^T(N + K)^{-1}(\mathrm{D} - \mathrm{DM}(t)) \\ &\qquad -\frac{1}{2}\ln\det(N + K) - \frac{n}{2}\ln(2\pi), \end{aligned} \tag{10}$$

where N is the noise matrix of the data and n is the dimension of N. The likelihood was given a uniform prior and the posterior was then obtained via MCMC sampling [1].

We allowed our Gaussian process to have knowledge of the data up to an arbitrary point and then let it attempt to predict future trends. Assuming our model is sufficient, we can use this technique to identify outliers in new data that may be indicative of a DM event.

During our runs, we allowed the Gaussian process to have access to data up to a specified date and then allowed it to make predictions of trends in the remainder of our existing data. We then sampled a fraction of the Gaussian fits and calculated the significance of each data point with the formula

$$\sigma = \left| \frac{\mathrm{D} - \mu_{\mathrm{Gauss}}}{\sqrt{\sigma_{\mathrm{D}}^2 + \sigma_{\mathrm{Gauss}}^2}} \right|, \tag{11}$$

where μ_{Gauss} is the mean of the sampled Gaussian fits, D is the data, σ_{D} is the standard deviation of the data, and σ_{Gauss} is the standard deviation of the Gaussian fits.

3 Results

3.1 Simple Baysian Analysis

For each pulsar modeled with this method, we present one figure containing three plots, with the first containing the model's fit over the DMX values, the second showing the residuals from the model and regions of 1σ and 2σ variation from the mean, and the third displaying the calculated σ value for each observation of the pulsar.

PSR J1741+1351 (Figure 1) showed a consistent linear variation and generally remained well-behaved. Overall, we found that our modelsmanage to provide accurate predictions for all pulsars showing just a linear trend.

PSR B1855+09 (Figure 2) followed a fairly linear trend in the 9 year data set, but took on new behavior in the 11 year data set. As a result, our predictions diverged from the observed 11 year data. There were a few similar cases in which our models could detect significant changes in the overall trends, and our model predictions were rendered useless unless they were either trained on more data or had some knowledge of the deviating trend.

PSR J1918-0642 hinted at variations with higher order terms in addition to the annual variations present in most of the pulsars. We found that our model was consistent in the tracking of this pulsars DMX variations and gave satisfactory predictions for the periodicity in the 11 year data.

Our final example, which we discuss in greater detail in section 4, is PSR J1713+0747, which experienced a DM event during the time of the 9 year dataset. The high sigma value generated for this pulsars DM event led us to believe that we can resolve significant variations in future DMX observations.

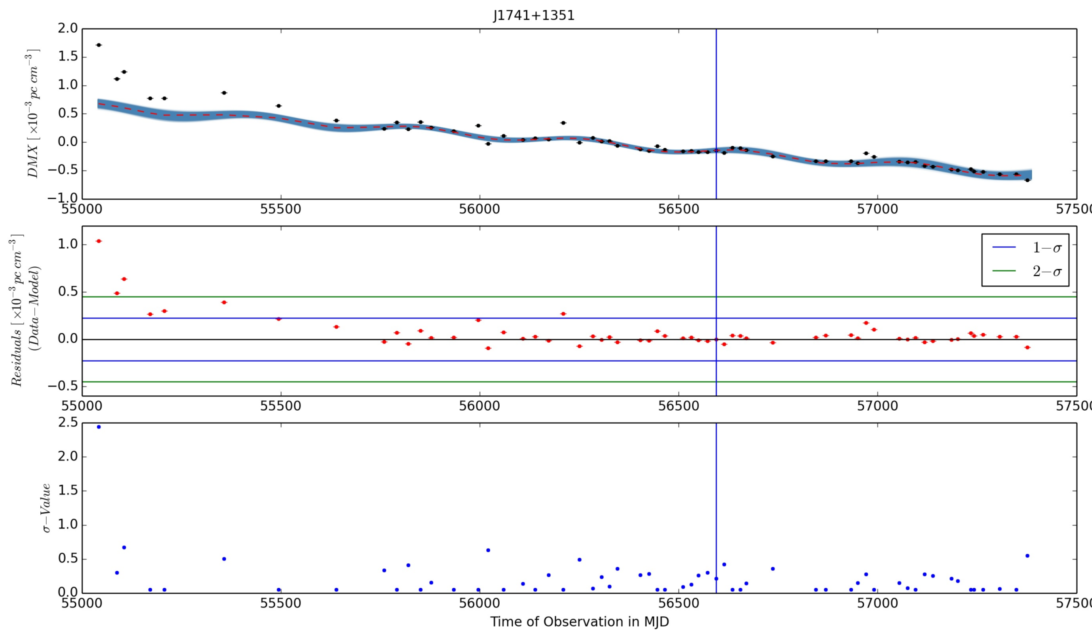

Figure 1: PSR J1741+1351 shows a linear trend in DMX variations. Our models from the 9 year dataset give accurate predictions for the 11 year period. The blue vertical line indicates the end of 9 year dataset.

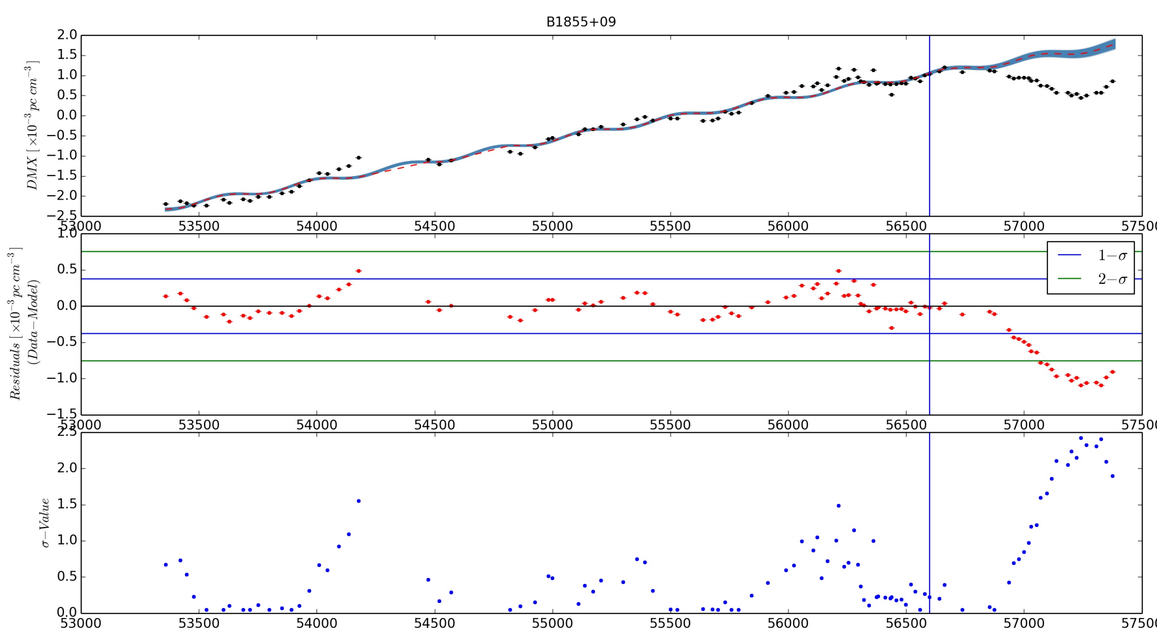

Figure 2: PSR B1855+09 shows a changing trend in its DMX variations. Our models detect this change in the overall trend of the pulsar but gives inaccurate predictions consequently. The vertical blue line indicates the end of the 9 year dataset.

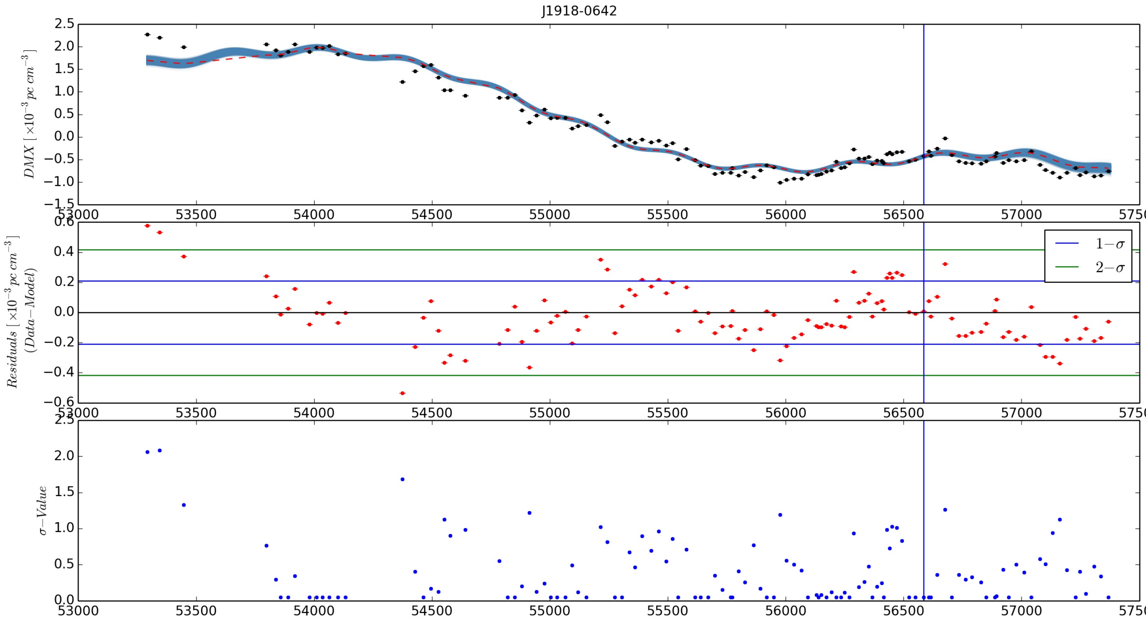

Figure 3: PSR J1918-0642 requires two sine terms and a linear term to model its variations. The blue vertical line indicates the end of the 9 year dataset.

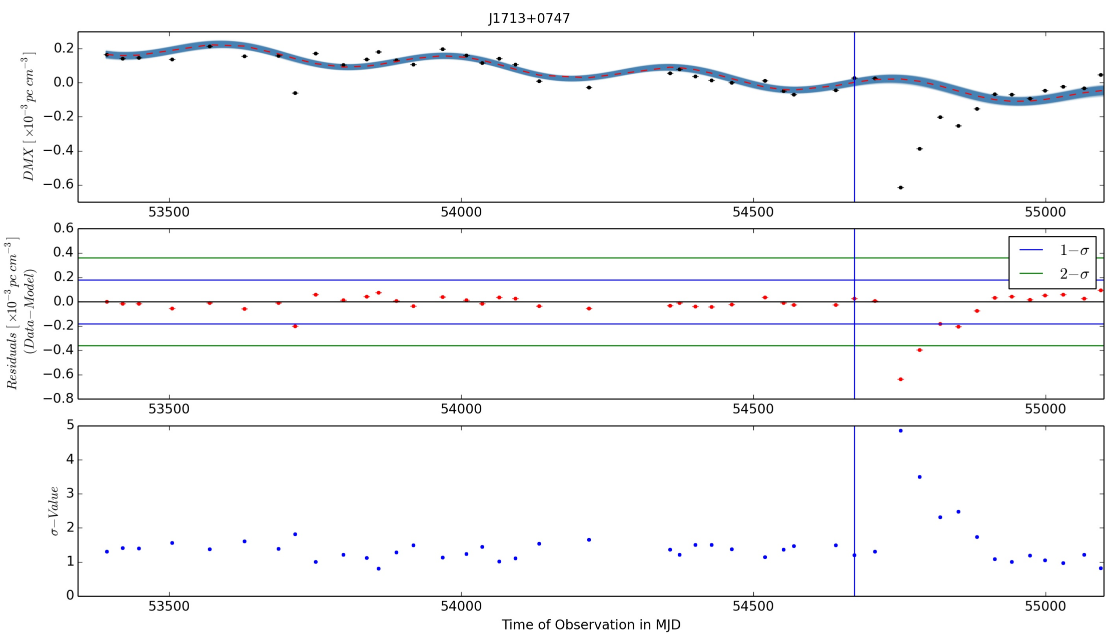

Figure 4: PSR J1713+0747 shows a DM event in the 9 year dataset. Our models can easily detect such events with high σ values. The blue vertical line indicates the date up to which data has been used to generate models

3.2 Gaussian Process Regression

As demonstrated by the fits in Figures 5 - 7, the Gaussian process method was exceedingly effective at modeling "visible" data. The method was also quite accurate at predicting trends up to a year in future for most variations, with results generally seeing a decline in accuracy after about 2 years for the strongest fits. A higher cadence and longer data sets would likely increase these time scales.

A significant result of this time window of accuracy is that significant deviations in DM trends (Figure 8), as well as DM events (Figure 9), might be found even with a yearlong gap in the data, assuming a relatively well-behaved ISM. Incidentally, significant changes in the ISM over many epochs may play a significant role in long-term predictability.

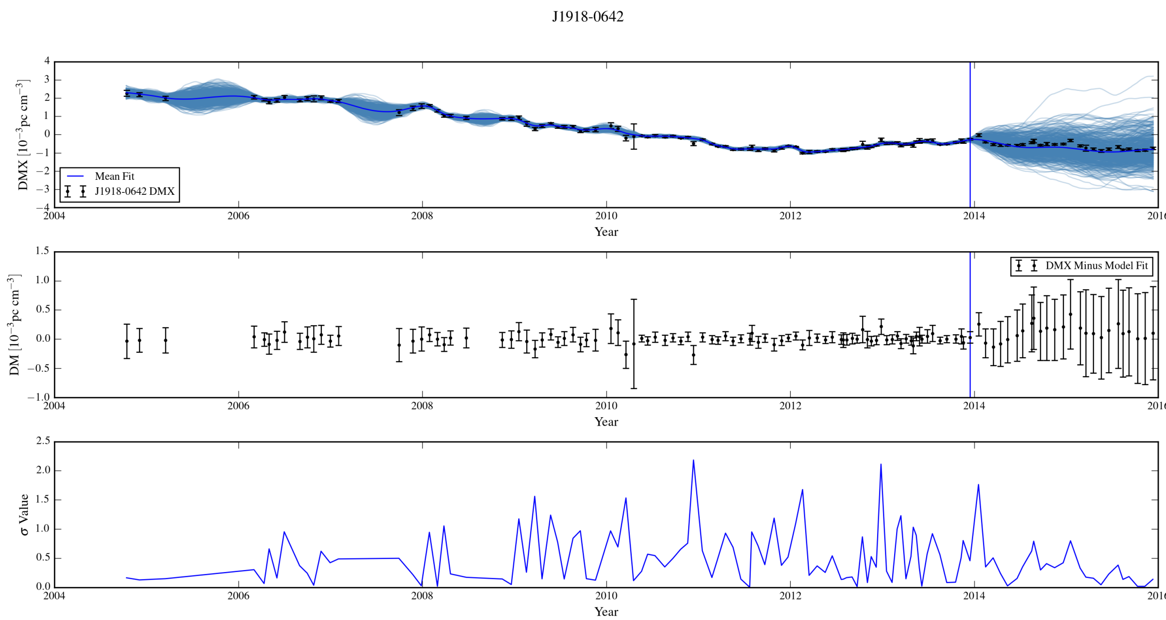

Figure 5: A Gaussian process on a pulsar with a strong linear DM trend. The vertical blue line indicates the point where the Gaussian process starts making predictions.

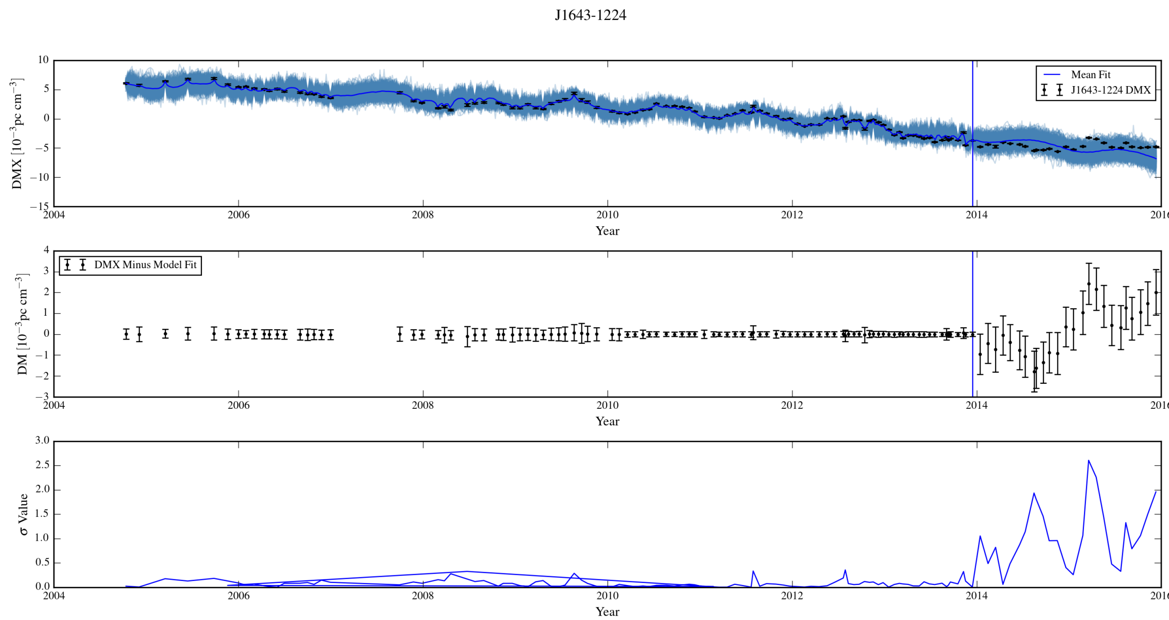

Figure 6: A Gaussian process on a pulsar with a strong periodic DM trend. The vertical blue line indicates the point where the Gaussian process begins making predictions.

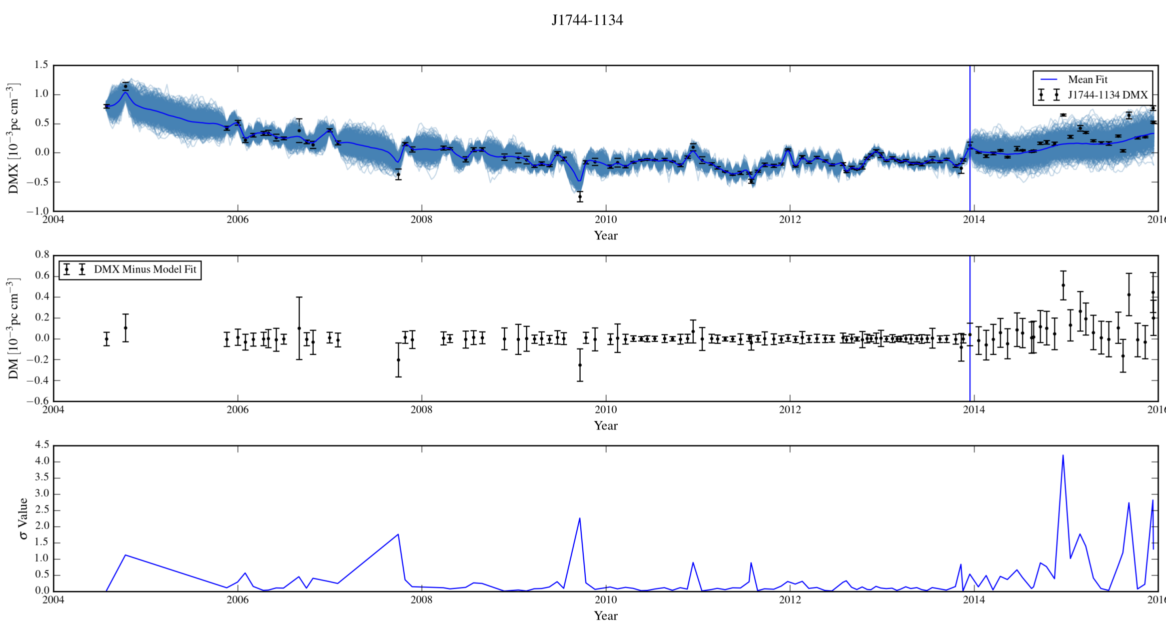

[!ht]

Figure 7: A Gaussian process on a pulsar with a strong quadratic DM trend. The vertical blue line indicates the point where the Gaussian process begins making predictions.

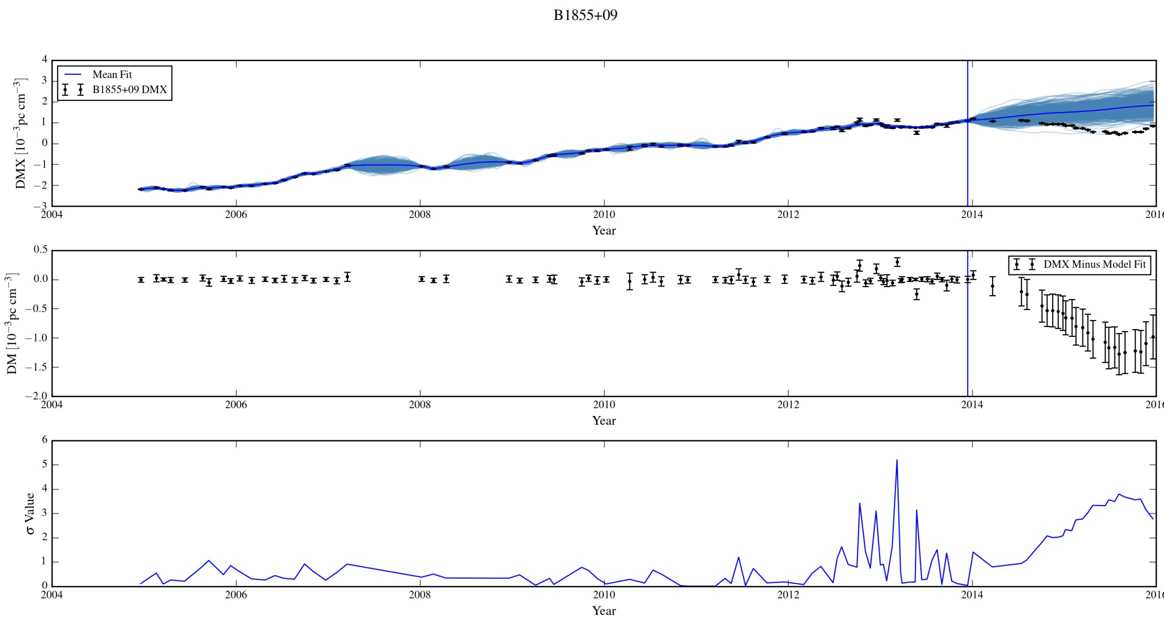

Figure 8: A Gaussian process on a pulsar that demonstrates a significant variation in its DM trend after 2014. The vertical blue line indicates the point where the Gaussian process begins making predictions.

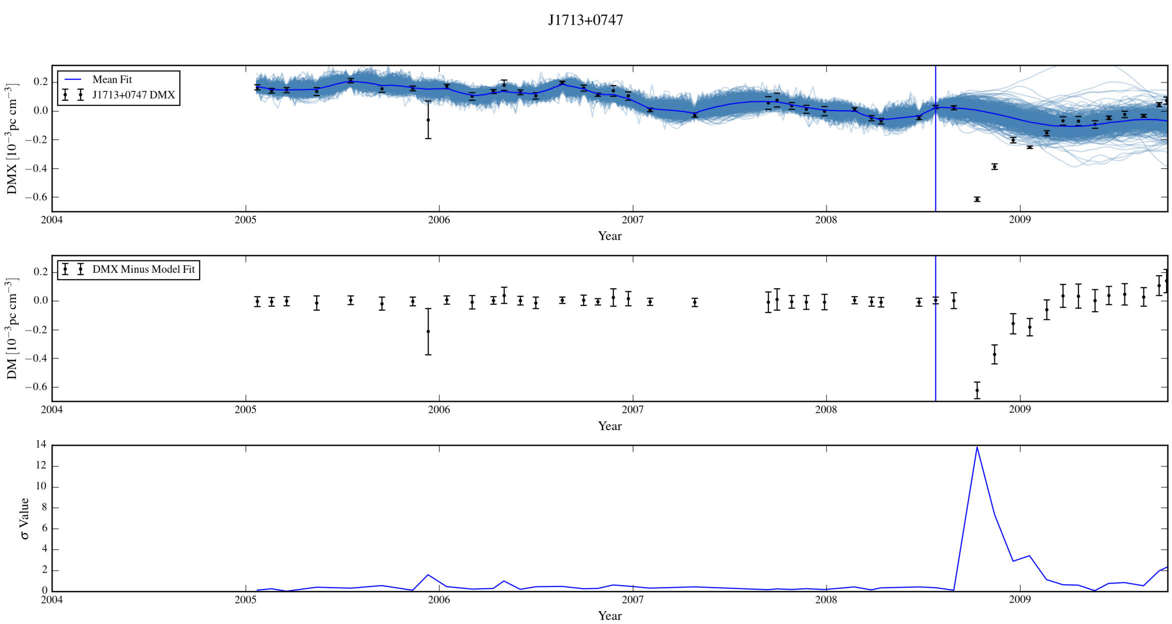

Figure 9: A Gaussian process recovering the DM event in PSR J1713+0747. The vertical blue line indicates the point where the Gaussian process beings making predictions.

4 The DM Event in PSR J1713+0747

We now turn to determining the physical characteristics of a DM event as seen in PSR J1713+0747.It is known that this particular pulsar is about 1.18 kpc away from Earth, moves with a proper motion of around 6.285 mas/yr and has a total DM_0 of approximately 15.99 pc cm^{-3} along the LOS[2] [11]. In Figure 10 we show a closeup of the sharp drop in the DMX variations observed around MJD 54751. From this figure we can make the following observations: 1.) The DM event is asymmetric with a sharp drop and gradual recovery. 2.) The DM event is aperiodic/unique as we have only seen it once. 3.) The sharpest decrease corresponds to a ΔDM of about -6×10^{-4}pc cm^{-3}. 4.) The DM recovery time is about 6 months.

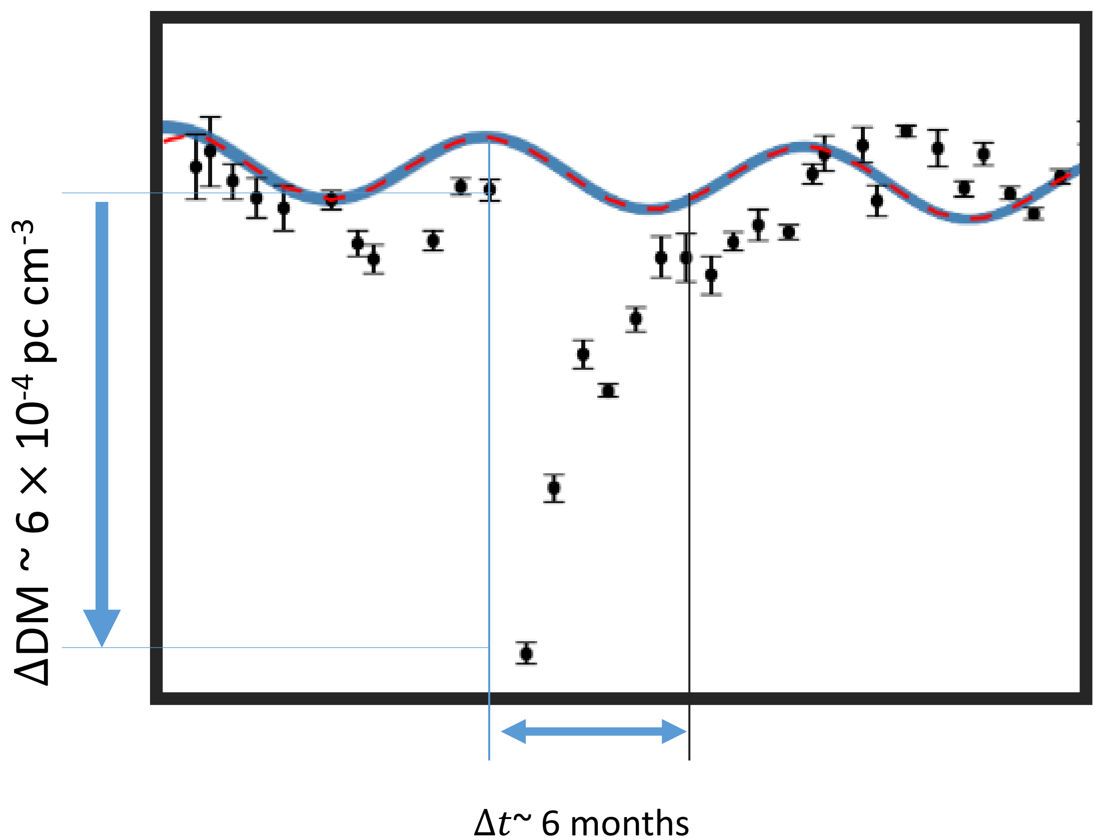

Figure 10: Closeup of the DM event in PSR J1713+0747

[2] http://www.atnf.csiro.au/people/pulsar/psrcat/

These observations lead us to conclude that our LOS must have crossed a region devoid of free electrons (which for the remainder of this paper will be referred to as a hole) in the ISM. For our analysis we made a couple of assumptions about the ISM in order to determine the properties of such a hole.

By rearranging equation 1, we found that a DM_0 of 15.99 pc cm^{-3} corresponded to a total of 10^{19} e^- along the LOS. Assuming that the ISM is homogenous, a drop of -6×10^{-4} pc implies that the hole must have been at least $\sim$ 10,000 AU in length along the LOS. Since we know that the hole moves across the line of sight in 6 months, if we assume the transverse velocity of the ISM to be negligible relative to transverse velocity of the Pulsar, we deduce that the hole needed to have an angular diameter of 3.1 mas. Additionally, if we assume that the hole is approximately halfway between Earth and the pulsar, the breadth of the hole across the line of sight should have been around 1.6 A.U.

These dimensional limits gave us some idea about the possible geometric structure of the hole. We considered many different shapes and settled on two basic toy models, in which the event could be described by either a cylinder or a crescent (or a half-crescent). In the cylinder model, we required the structure to have a negative density gradient along its length and that it be tilted at a small angle from the LOS. We determined that the hole should have an extremely low density of electrons along the initial LOS which increased gradually until it matched the density of electrons in the surrounding ISM. In the crescent model, the density remained constant while the length of intercept made by the LOS through the hole decreased gradually with time due to the curved shape. Both full and half-crescents satisfy the shape required for a gradually vanishing hole.

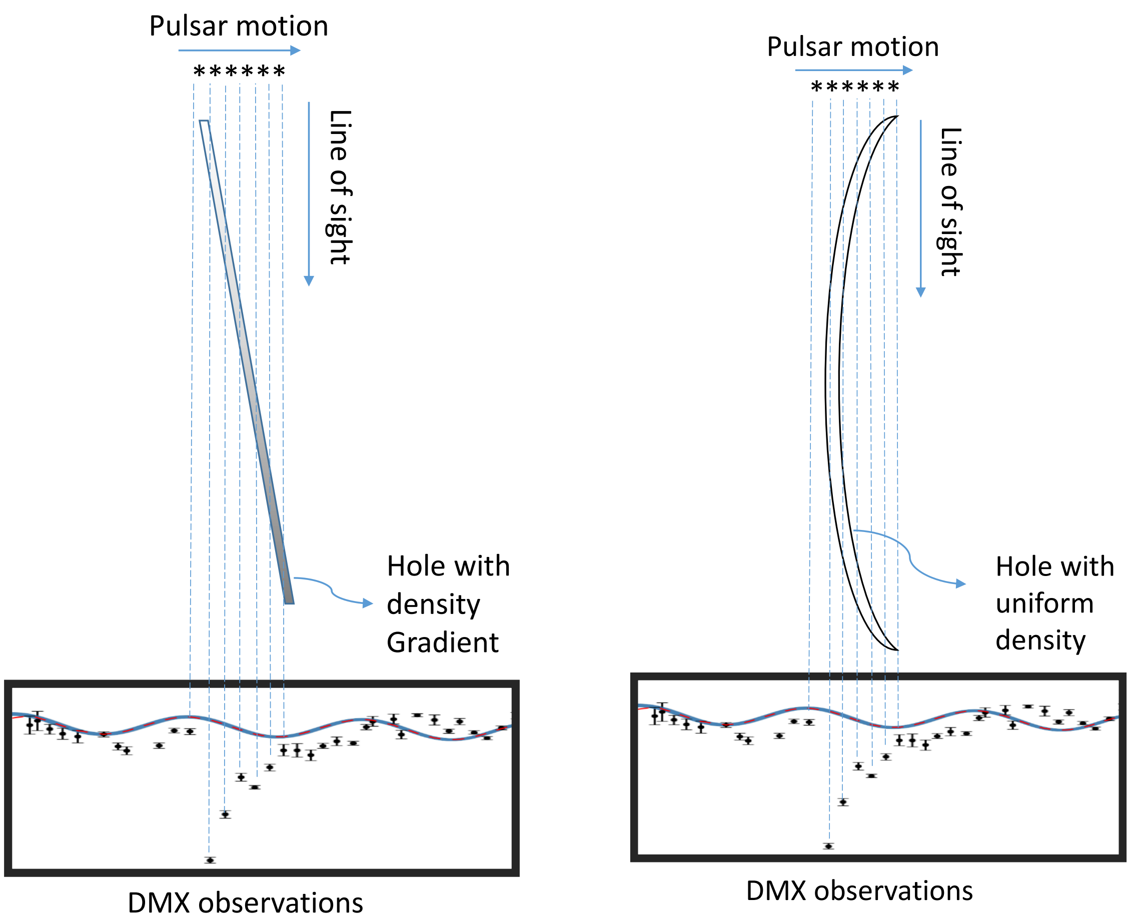

Figure 11: Model candidates for the DM event in PSR J1713+0747. (*Left*) Cylinder model. (*Right*) Cresent model.

We also attempted to identify possible physical systems which could cause a decrease in the electron density of the ISM. Based on the size and geometry of the structure, we postulate that the most likely candidates include interstellar filaments, magnetospheres of stars, and stellar wind shock waves. Each of these mechanisms has the ability to create regions in the ISM which are devoid of free electrons, as well as fit the proposed shapes. Another possiblity is that some A.U.-scale "shield" or enveloping layer of HII or a non-ionizing material such as dust could have protected this region of the ISM from ionization. This particular explanation fits well with the model (Figure 11) and is further strengthened by observations of dust and HI in regions near our LOS to the pulsar.

We can further narrow down our list by imposing additional physical constraints on the system, such as calculating a rough scale of the energy required to create the hole. However, it is likely that evidence in the form of A.U.-resolution images of the LOS to PSR J1713+0747 will be required to confidently rule out or confirm any of these candidates.

Given the highly disproportionate ratio of the structure's dimensions, it is also conceivable that, rather than observing the entire structure, we are instead observing a tiny sliver of a much larger architecture in the ISM. Such a configuration would likely be on the order tens of thousands of A.U. to a few pc in length and around 0.1 pc in width [2]. Indeed, in H-α, HI, and dust images near the pulsar, we observe structures of this magnitude, making it a possibility that some filament branching off from such a region could have crossed our LOS.

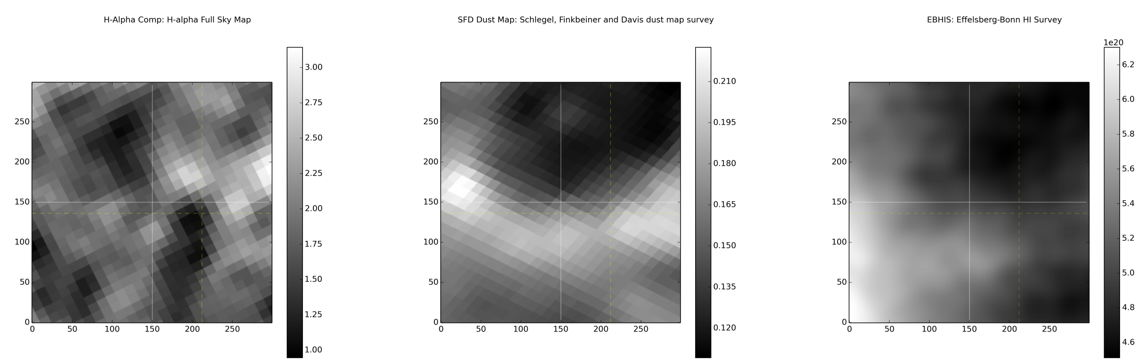

Figure 12: Observations at various frequencies in the direction of PSR J1713+0747. Images are 1 deg^2 in area, with the x- and y axes labeled in pixels and the colorbars in arbitrary units. White cross indicates the location of the pulsar in the image. (*Left*) Hα (from [5]); (*Middle*) Dust distribution, as inferred from the Schlegel, Finkbeiner, & Davis Dust Map Survey [13]; and (*Right*) H I, from the Effelsberg-Bonn Survey [15]

5 Future Work

A better understanding of DM events requires their real-time detection. As of this writing, we are in the process of implementing a near real-time system to process TOAs with PSRCHIVE[3] and analyze them for DM

[3] psrchive.sourceforge.net

variations with TEMPO[4] and TEMPO2[5]. In doing so, we will be able to catch significant variations as soon as they appear in the data.

We can further narrow down our search for the physical structures behind these events by imposing additional constraints on the system, such as calculating a rough scale of the energy required to create the hole. However, it is likely that evidence in the form of A.U.-resolution images of the LOS to PSR J1713+0747 will be required to confidently rule out or confirm any of these candidates. As such, it would prove useful to have access to observatories that offer A.U.-level resolution on an as-needed basis.

6 Summary & Conclusions

We have analyzed DM variations in the pulsars of the NANOGrav collaboration in an attempt to effectively predict future trends as well as identify significant deviations caused by structures in the ISM. Simple Bayesian and Gaussian process regression methods were used to generate models and predict these variations and have been found to work effectively up to two years into the future, in most cases.

We also analyzed the DM event found in the data of PSR J1713+0747 and concluded that such an event must be caused by a region in the ISM devoid of electrons that measures approximately 1.64 au in width transverse to the LOS and 10^4 au along the LOS. Based on the characteristics of the deviation, we proposed a few toy models and considered the astrophysical mechanisms for their creation. Finally, we discussed the ongoing implementation of a program to analyze TOAs and analyze DM variations in near real-time. We also concluded that access to telescopes with au-scale resolution would be necessary to further improve our understanding of DM events.

We thank J. Ellis, S. Taylor, and M. Vallisneri for their guidance on Bayesian analysis. We thank the team at the SFP and ISP Offices at Caltech for ensuring that our logistics were taken care of. Finally, we thank our financial sponsors at SURF 2016, NSF, Oberlin College, and IIT(BHU) Varanasi, for providing us with the funnding to make our stays at Caltech possible. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The NANOGrav project receives support from NSF Physics Frontier Center award number 1430284.

[4] tempo.sourceforge.net
[5] https://bitbucket.org/mkeith/tempo2

Appendix A

Simple Bayesian Analysis Code Summary

1. Read in file containing new data
2. Read in the previous data-set file
3. Read in guesses
4. Count the number of sine terms required by counting the number of guesses given
5. Try to optimize the guesses using scipy.optimize
6. Set-up the MCMC sampler
7. Define log prior function
8. Define log likelihood function
9. Define a function to generate models
10. Define log posterior function
11. Run a 100 step burn-in
12. Reset sampler
13. Run the sampler again for large number of steps, output saved in chains
14. Burn-in the first 25% of the resulting chains
15. Reshape the chains into flatchains
16. Generate models and predictions using flatchains
17. Calculate mean of models and predictions upto the new data point
18. Generate an array of gaussians (with mean = data value and std. dev. = errorbar) representing the array of data points
19. Calculate array of residual distributions by subtracting the distribution of models from the array of data gausssians.
20. Calculate the mean and standard deviation of each individual residual in the residual distribution array.
21. Calculate the mean and standard deviation of the whole residual array.
22. Similarly calculate residual distribution for the new data point by subtracting the predicted distribution from a gaussian generated for the new data point.

23. Calculate mean and standard deviation of residual distribution of the new data point. (This mean is taken as the residual value for the new data point and the standard deviation as the error in calculating the residual)

24. Calculate $\sigma -$ value for new data point as :

$$\frac{\text{(Residual of new data point - Mean of residuals of previous dataset) - Error in calculating the residual of new data point}}{\text{Standard deviation of residuals of previous dataset}}$$

25. Print results.

26. Compare the $\sigma -$ value with a threshold, and rasie an alarm if found larger.

Gaussian Process Regression Code Summary

1. Read in the current data set file.

2. Read in file containing new data.

3. Construct a model with no quadratic term.

4. Add random Gaussian variables to the model.

5. Define the according log prior and log likelihood functions.

6. If the log Bayes factor is already known, ignore step 3 and run the PTMCMC sampler with the proper model. Otherwise, run the PyMultinest sampler to obtain the log evidence.

7. If PyMultinest was used in the previous step, repeat all above steps for a model with the quadratic term included. Otherwise, ignore this step.

8. If PyMultinest was used, compute the log Bayes factor from the log evidences of the two models to determine which model to use.

9. Draw 500 samples and use them to predict trends in the data after a certain epoch.

10. Compute the mean and standard deviation of those samples.

11. Plot the posterior distributions, as well as a triple plot of the data with the 500 model samples and mean, the residuals, and the corresponding significance values at each point.

References

[1] S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O'Neil. Fast Direct Methods for Gaussian Processes and the Analysis of NASA Kepler Mission Data. March 2014. 3

[2] D. Arzoumanian, P. André, N. Peretto, V. Könyves, N. Schneider, P. Didelon, and P. Palmeirim. Characterizing interstellar filaments with Herschel in nearby molecular clouds. In *From Atoms to Pebbles: Herschel's view of Star and Planet Formation*, March 2012. 13

[3] Buchner, J., Georgakakis, A., Nandra, K., Hsu, L., Rangel, C., Brightman, M., Merloni, A., Salvato, M., Donley, J., and Kocevski, D. X-ray spectral modelling of the agn obscuring region in the cdfs: Bayesian model selection and catalogue. *A&A*, 564:A125, 2014. 3

[4] The NANOGrav Collaboration, Zaven Arzoumanian, Adam Brazier, Sarah Burke-Spolaor, Sydney Chamberlin, Shami Chatterjee, Brian Christy, James M. Cordes, Neil Cornish, Kathryn Crowter, Paul B. Demorest, Timothy Dolch, Justin A. Ellis, Robert D. Ferdman, Emmanuel Fonseca, Nathan Garver-Daniels, Marjorie E. Gonzalez, Fredrick A. Jenet, Glenn Jones, Megan L. Jones, Victoria M. Kaspi, Michael Koop, Michael T. Lam, T. Joseph W. Lazio, Lina Levin, Andrea N. Lommen, Duncan R. Lorimer, Jing Luo, Ryan S. Lynch, Dustin Madison, Maura A. McLaughlin, Sean T. McWilliams, David J. Nice, Nipuni Palliyaguru, Timothy T. Pennucci, Scott M. Ransom, Xavier Siemens, Ingrid H. Stairs, Daniel R. Stinebring, Kevin Stovall, Joseph K. Swiggum, Michele Vallisneri, Rutger van Haasteren, Yan Wang, and Weiwei Zhu. The nanograv nine-year data set: Observations, arrival time measurements, and analysis of 37 millisecond pulsars. *The Astrophysical Journal*, 813(1):65, 2015. 1

[5] D. P. Finkbeiner. A Full-Sky Hα Template for Microwave Foreground Prediction. *The Astrophysical Journal Supplement Series*, 146:407–415, June 2003. 13

[6] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *Publ. Astron. Soc. Pac.*, 125:306–312, November 2013. 3

[7] M. P. Hobson, A. H. Jaffe, A. R. Liddle, P. Mukherjee, and D. Parkinson. *Bayesian Methods in Cosmology*. February 2014. 2

[8] Robert E Kass and Adrian E Raftery. Bayes factors. *Journal of the american statistical association*, 90(430):773–795, 1995. 3

[9] M. T. Lam, J. M. Cordes, S. Chatterjee, M. L. Jones, M. A. McLaughlin, and J. W. Armstrong. Systematic and stochastic variations in pulsar dispersion measures. *The Astrophysical Journal*, 821(1):66, 2016. 1, 2

[10] T. B. Littenberg and N. J. Cornish. Bayesian approach to the detection problem in gravitational wave astronomy. *Phys. Rev. D.*, 80, 2009. 3

[11] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. The Australia Telescope National Facility Pulsar Catalogue. *The Astrophysical Journal*, 129:1993–2006, April 2005. 10

[12] R. Ramachandran, P. Demorest, D. C. Backer, I. Cognard, and A. Lommen. Interstellar plasma weather effects in long-term multifrequency timing of pulsar b1937+21. *The Astrophysical Journal*, 645(1):303, 2006. 2

[13] D. J. Schlegel, D. P. Finkbeiner, and M. Davis. Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds. *The Astrophysical Journal*, 500:525–553, June 1998. 13

[14] Rutger van Haasteren and Michele Vallisneri. New advances in the gaussian-process approach to pulsar-timing data analysis. *Phys. Rev. D*, 90:104012, Nov 2014. 3

[15] B. Winkel, J. Kerp, L. Flöer, P. M. W. Kalberla, N. Ben Bekhti, R. Keller, and D. Lenz. The Effelsberg-Bonn H i Survey: Milky Way gas. First data release. *Astronomy & Astrophysics*, 585:A41, January 2016. 13