TempoNest tutorial - IPTA student workshop, 2018

Aditya Parthasarathy and Ryan Shannon

June 14, 2018

Abstract

I hope that we all know and appreciate by now, the usefulness of precision pulsar timing. The exceptional rotational stability of millisecond pulsars combined with sophisticated and powerful analysis tools has led to the pursuit of a variety of science goals, from the detection of planets around pulsars to some of the most precise tests of general relativity in the strong field regime.
In this tutorial you will be briefly learning how to use TempoNest for modelling timing noise in pulsars.
TempoNest is a Bayesian analysis software that performs a **simultaneous analysis** of the linear or non-linear deterministic pulsar timing model and additional stochastic parameters. It uses the Bayesian inference tool *MultiNest* to explore this joint parameter space, whilst using Tempo2 as an established means of evaluating the timing model at each point in that space.
Please refer to **Lentati et al. (2014)** and *https://github.com/LindleyLentati/TempoNest* for more detailed information.

1 General Introduction

Let's start with a brief recap of pulsar timing. The arrival times of pulses (ToAs) from a pulsar are recorded by an observatory in a series of discrete observations over a certain period time. To account for the motion of the Earth, these ToAs must be transformed into a common frame of reference, the solar system barycenter. Once that is done, a model of the pulsar is fit to the ToAs. This model generally describes the position, spin down, orbital parameters (if it's a binary) etc. Any discrepancies between the observed ToAs and those predicted by the pulsar model are manifested in the timing residuals. The timing residuals are extensively analyzed to understand and model the observed discrepancies. For a more detailed review of pulsar timing, start by referring to **Lorimer & Kramer (2004)**.

Tempo2 is the most commonly used software for the above analysis. While fitting the pulsar model to the measured ToAs, it uses an initial guess to generate pre-fit timing residuals and then a linear least squares fit is used to improve the timing residuals. Typically multiple iterations of the fit are performed until convergence is reached.

2 What's the problem then?

The assumption made in Tempo2 is that the ToAs can be solely represented as a sum of the deterministic signal due to the timing model and uses a white noise model that is dictated only by the uncertainties of the ToAs.
The issue with this assumption is that, it is **not true** in most realistic data sets.

Power from stochastic signals like timing noise, which is predicted to be due to the rotational irregularities of the neutron star or the correlated noise due to the stochastic gravitational wave backgound (GWB) signal, is absorbed by the timing model, hence affecting the process of parameter estimation and also breaking down the assumptions made previously.

Over the years, there have been several techniques developed to account for this issue. The Cholesky decomposition method as described in **Coles et al. (2011)** uses the Cholesky decomposition of the covariance matrix that is estimated from the power spectral density of the timing residuals to whiten the residuals. However as stated later by **van Haasteren & Levin (2013)**, the post fit timing residuals aren't time stationary making the estimation of the power spectral density poorly defined.

So, perhaps the preferred approach is to perform a joint analysis of the timing parameters simultaneously along with the stochastic parameters.

The above description of the problems and techniques is by no means a comprehensive one. Timing noise in pulsars, young and old, have been studied over many years (as you would have seen from Ryan's talk) and providing a descriptive review is beyond the scope of this tutorial.

3 What does TempoNest do?

TempoNest provides a means of performing a simultaneous analysis of linear (or) non-linear timing model parameters with stochastic parameters using a Bayesian inference tool called **MultiNest** to efficiently sample this joint parameter space. It also uses Tempo2 to evaluate the solution at each point in this space and finally provides a Bayesian evidence value to compare and select different models used for the modelling the data.

Now, if you didn't follow through the last paragraph, these are key words that you might need to focus on, for the purposes of this tutorial:

- Bayesian inference
- Bayesian evidence
- Model selection

3.1 What's Bayesian inference?

If we have a set of model parameters defined as Θ in a model or hypothesis, H given the data, D. Bayes' theorem states that:

$$\Pr(\Theta \mid D, H) = \frac{\Pr(D \mid \Theta, H)\Pr(\Theta \mid H)}{\Pr(D \mid H)}, \tag{1}$$

where,

- $\Pr(\Theta \mid D, H) \equiv \Pr(\Theta)$ is the posterior probability distribution of the parameters,
- $\Pr(D \mid \Theta, H) \equiv L(\Theta)$ is the likelihood,
- $\Pr(\Theta \mid H) \equiv \pi(\Theta)$ is the prior probability distribution,
- and $\Pr(D \mid H) \equiv Z$ is the Bayesian Evidence.

3.2 What's Bayesian evidence?

From the above equation, we can estimate the evidence by the following equation:

$$Z = \int L(\Theta)\pi(\Theta)\mathrm{d}^n\Theta, \tag{2}$$

where n is the dimensionality of the parameter space.

3.3 What's model selection?

The equation for evidence **Z** states that it is the average of the likelyhood over the prior. That states that:

- The evidence is larger for a model if more of its parameter space is likely
- and smalled for a model if its parameter space has low likelyhood values

Thus the evidence is a very useful parameter that helps us compare and select the best model that is significantly better at explaining the data.

4 How does TempoNest work?

Briefly, TempoNest uses the Bayesian inference package called MultiNest, which is a nested sampling framework which provides an efficient means of sampling from posteriors and also calculates the evidence. Please refer to the **Feroz et al. (2009); Cameron & Pettitt (2013)** for a more detailed understanding of the algorithm.

5 How does all of this apply to pulsar timing?

For any given pulsar, we can describe the ToAs for the pulses as a sum of a deterministic and a stochastic component.

$$\mathbf{t}_{\text{total}} = \mathbf{t}_{\text{deterministic}} + \mathbf{t}_{\text{stochastic}}, \tag{3}$$

Here is an excerpt from the **temponest paper** that is quite important and perhaps summarizes the model used for the analysis of data.

Writing the deterministic signal due to the timing model as $\tau(\epsilon)$, and the uncertainty associated with a particular TOA i as σ_i we can write the likelihood that the data is described solely by the timing model as:

$$\Pr(\mathbf{t}|\epsilon) \propto \left(\prod_{i=1}^{n} \sigma_i^2\right)^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\sum_{i=1}^{n}\frac{(t_i - \tau(\epsilon)_i)^2}{\sigma_i^2}\right). \quad (4)$$

This represents the simplest model choice possible in TempoNest, including only those free parameters present in the TEMPO2 fit. From here we can now begin to ***make our model for the stochastic contribution to the signal more realistic by introducing additional parameters to describe the white and red noise components****, in order to compare the evidence with this simpler model and determine the optimal set of parameters supported by the data.*

If you're interested in understanding how the white noise and red noise processes are modelled and how TempoNest uses all of this to marginalize over the timing model, please refer to the **Lentati et al. (2014)**

We will, however cover some of the noise modelling aspects in the following sections.

5.1 Using TempoNest: a brief introduction

For this exercise, you will need to download a few files from a github repository. The details will be announced. Make sure that you have the following files before proceeding:

- J1539-5626.par
- J1539-5626.tim
- A directory named *config_files* containing PM_steepRN.cfile, propermotion.cfile and red_noise.cfile.

In this exercise we shall be trying to model the timing noise in a young pulsar, J1539-5626. Try using *psrcat* to find out more about this pulsar.

You might be familiar at this stage with the contents of the .par and .tim files. Use *tempo2* and plot the timing residuals. How do they look like? What can you find out about the data set?

You will notice that there seems to be a strong presence of red noise in the timing residuals. Young pulsars are known to exhibit significant timing noise, perhaps due to rotational irregularities of the neutron star. We will try using TempoNest to model and subtract the red noise from the timing residuals.

Since TempoNest is built as a plugin for Tempo2, you can run TempoNest using:

```
tempo2 -gr temponest -f <par> <tim> -cfile <config>
```

But before we go ahead and run temponest, we have to understand a few things.

5.2 The TempoNest configuration file

Let's try to understand the configuration file for TempoNest. Open up any of the **.cfile** from the *config_files* directory.

The parameters set in the configuration file dictates the settings used in that TempoNest run.

For a detailed explanation of the configuration file, please refer to the TempoNest manual that you must have pulled along with the other files from the github repository. Please note that this is an outdated manual but for the purposes of this exercise, it's still valid.

5.2.1 Key parameters in the configuration file for this tutorial

What do you think are the key differences between the different configuration files?

You must have noticed that the **root** parameter is different among the configuration files. This parameter specifies the output directory for the TempoNest results and the prefix for the filenames. For example,

```
root = <directory>/<prefix>-
```

Please note that it is required for you to create the output directory manually and TempoNest does not do this for you.

The other key difference is that, the **custom priors** flag is set to 1 in a couple of cases and to 0 for one. This parameter, when True (set to 1) allows TempoNest to fit for specific parameters as specified using **TempoPriors**.

The index number specified in **TempoPriors[index][2]=0** corresponds to the index number of the fit parameter in the .par file. So, if *F0* is the third fit parameter in the .par file, then to enable TempoNest to use this parameter as a custom prior, we set *customPriors=1* and set *TempoPriors[3][2]=1*.

Finally, the last key difference is that one of the config files has an additional flag, **FitLowFreqCutoff=1**. This allows TempoNest to fit for a steeper Red noise model.

Though the prior ranges are the same across the different configuration files, it is important to note that the parameters **EFACPrior, EQUADPrior,AlphaPrior,AmpPrior** set the prior probability distributions for the EFAC, EQUAD, Red noise slope and Red noise amplitude.

Please refer to the TempoNest manual or "ask questions" if you wish to understand the other parameters present in the configuration file.

5.3 Running TempoNest

Try running TempoNest using the command mentioned above using the 3 different configuration files. Please remember to create the required output directories.

You'll notice that a lot of "stuff" is printed out during this process and it stops once the fit converges after multiple iterations. This might take a few minutes depending on the configuration file, but be patient!

Read through the output and see if it any of this makes sense, if not, either refer to the manual or ask people!

5.4 TempoNest results

Let us finally look at the results that TempoNest produces. You'll notice that TempoNest has produced a series of files with the prefix as assigned in the configuration file. The key files that we might have to look at for this tutorial are the following:

- The output par file
- The paramnames file
- The post_equal_weights.dat file
- The stats file

TempoNest outputs a new .par file after each run. What difference do you notice between the new par file and the old one (the input par file)?

There are a set of new parameters at the end of the new par file output by TempoNest. These parameters, **TNGlobalEF,TNGlobalEQ, TNRedAmp and TNRedGam** describe the white and red noise models. So what did TempoNest do? To take a look at the result, use tempo2 to load the new par file with the old tim file.

```
tempo2 -gr plk -f <new_par> <old_tim>
```

Look at the pre-fit and the post-fit residuals. What do you notice? Now look at the Red-noise model? Does this look similar to the structure in the post-fit residuals? Now press *'shift+k'*. What do you notice? Does the noise modelling work as expected?

Try the same for all of the new par files. What difference do you notice among the different models?

Now, let's take a look at the posterior distributions. View the post_equal_weights.dat file. The columns in this file are in the same order as mention in the .paramnames file. Try plotting the different columns. What do you see? This file contains the equally weighted posterior samples. The posteriors are important in validating the results and are also important to study the correlation among the different fit parameters. What do you understand from the posterior distributions from the different TempoNest runs?

Finally, let's look at comparing the evidence of the different models and how it can be used to select the best model. TempoNest provides an evidence value for each model and this can be seen in the stats.dat file. The evidence value that we want to look at is the **Nested Importance Sampling Global Log-Evidence**. Compare these values for the different models. Which do you think is the best model? Why do you think the other models have lower evidence values?

References

Cameron E., Pettitt T., 2013, preprint, https://ui.adsabs.harvard.edu/abs/2013arXiv1309.2737C p. arXiv:1309.2737 (http://arxiv.org/abs/1309.2737 `arXiv:1309.2737`)

Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, http://dx.doi.org/10.1111/j.1365-2966.2011.19505.x , https://ui.adsabs.harvard.edu/abs/2011MNRAS.418..561C 418, 561

Feroz F., Hobson M. P., Bridges M., 2009, http://dx.doi.org/10.1111/j.1365-2966.2009.14548.x , https://ui.adsabs.harvard.edu/abs/2009MNRAS.398.1601F 398, 1601

Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, http://dx.doi.org/10.1093/mnras/stt2122 , https://ui.adsabs.harvard.edu/abs/2014MNRAS.437.3004L 437, 3004

Lorimer D. R., Kramer M., 2004, Handbook of Pulsar Astronomy

van Haasteren R., Levin Y., 2013, http://dx.doi.org/10.1093/mnras/sts097 , https://ui.adsabs.harvard.edu/abs/2013MNRAS.428.1147V 428, 1147