

The Thousand-Pulsar-Array programme on MeerKAT XIX: single-pulse data analysis, nulling and pulse energy distributions

Michael J. Keith ^{1,★}, Patrick Weltevrede ¹, Lucy Oswald ^{2,3,4}, Aris Karastergiou,⁴ Xiaoxi Song ⁵, Haoyue Wang ¹, Jui-An Hsu ¹, Simon Johnston ⁶, Geoff Wright,¹ Matthew Bailes^{7,8} and Maciej Serylak⁹

¹Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK

²School of Physics & Astronomy, University of Southampton, Southampton SO17 1BJ, UK

³Astrophysics Group, Cavendish Laboratory, University of Cambridge, J. J. Thomson Avenue, Cambridge CB3 0HE, UK

⁴Department of Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK

⁵ASTRON, the Netherlands Institute for Radio Astronomy, Oude Hoogeveensedijk 4, NL-7991 PD Dwingeloo, the Netherlands

⁶Australia Telescope National Facility, CSIRO, Space and Astronomy, PO Box 76, Epping, NSW 1710, Australia

⁷Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, VIC 3122, Australia

⁸OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Hawthorn, VIC 3122, Australia

⁹SKA Observatory, Jodrell Bank, Lower Withington, Macclesfield SK11 9FT, UK

Accepted 2026 June 8. Received 2026 May 29; in original form 2026 March 30

ABSTRACT

We present the Thousand Pulsar Array (TPA) single-pulse data release, obtained with the MeerKAT radio telescope and comprising time-series observations of 1192 pulsars, typically containing ~ 1000 consecutive pulses per source. We describe the MeerTime Single Pulse software pipeline which calibrates the data and automatically excises interference signals to produce data products suitable for typical single-pulse studies. To demonstrate the capabilities of the data set, we carry out a population-level study of phase-averaged single-pulse energy distributions and nulling behaviour. Pulse energy distributions are modelled within a Bayesian framework choosing from a range of intrinsic energy distributions, and including an explicit nulling fraction. We find that approximately half of the pulsars require multicomponent intrinsic energy distributions, while the remainder are consistent with single-component models. Nulling is detected or constrained for most pulsars in the sample, and both the occurrence and inferred nulling fraction show systematic variation across the $P-\dot{P}$ diagram. In particular, nulling fractions increase with spin period and exhibit only a weak dependence on period derivative. We also examine trends in the preferred forms of pulse energy distributions as a function of spin-down luminosity, finding modest evidence for population-level evolution. Estimates of single-pulse luminosities indicate that individual pulses can exceed the long-term average luminosity by large factors, particularly for low- $\dot{E}$ pulsars. These results characterize the statistical properties of single-pulse emission across a large pulsar sample and highlight the limitations of phase-averaged energy distributions for capturing the full complexity of pulsar emission variability.

Key words: pulsars: general.

1 INTRODUCTION

Although the time-averaged shape of pulses from radio pulsars is usually highly stable, individual pulses can vary considerably from pulse to pulse. Pulse-to-pulse variability of pulsars comes in many forms, including nulling (D. C. Backer 1970), mode-changing (e.g. Z. G. Wen et al. 2020), variation in pulse intensity (e.g. F. G. Smith 1973), random phase jitter of subpulses (e.g. A. Parthasarathy et al. 2021; T. Kikunaga et al. 2024), or organized drifting or modulation of subpulses (e.g. P. Weltevrede, B. W. Stappers & R. T. Edwards 2007; X. Song et al. 2023). The time

variability from radio pulsars is important to investigate and understand, both because doing so directly probes the physics of the neutron star magnetosphere (e.g. X. Song et al. 2023), and because such variability presents a problem for precision pulsar timing for tests of gravitation and searches for gravitational waves (e.g. R. J. Jennings et al. 2024).

Previous studies of pulsar pulse energy distributions have suggested that typical pulsar emission energies are drawn from a lognormal distribution (e.g. I. H. Cairns, S. Johnston & P. Das 2001, 2004), although M. B. Mickaliger et al. (2018) suggested that a power law divided by an exponential leads to a better fit for most of their survey of 264 pulsars and 14 rotating radio transients, and showed that some pulsars deviate from a log-normal power distribution due to excess high-energy pulses. S. Burke-Spolaor et

* E-mail: mkeith@pulsarastronomy.net

al. (2012) found that, for 315 pulsars detected in the mid-latitude portion of the Parkes High Time Resolution Universe survey, about 40 per cent showed a lognormal single-pulse energy distribution and the rest showed behaviour that was characterized as either Gaussian or erratic. Separately, it seems that giant pulses tend to follow a power-law distribution (S. C. Lundgren et al. 1995; S. Johnston & R. W. Romani 2002), and it is suggested that giant pulses originate not from near the polar cap, like other radio pulses, but in or close to the current sheet instead (Y. Lyubarsky 2019; A. Philippov et al. 2019).

Nulling in pulsars is seen over a wide range of time-scales, with each null lasting typically fewer than five pulses (R. Basu, D. Mitra & G. I. Melikidze 2017), but long-term nulling can be seen where emission ceases for months or years (e.g. M. Kramer et al. 2006). Nulls are often seen to have a characteristic time-scale or quasi-periodic behaviour for a given pulsar, which is typically different to that of the subpulse modulation or drifting, indicating a different origin to the nulling (R. Basu et al. 2017). Nulling is observed to be a broad-band phenomenon (P. Weltevrede, R. T. Edwards & B. W. Stappers 2006; P. Weltevrede et al. 2007; V. Gajjar et al. 2014; A. Naidu et al. 2017) but can also appear only in selective frequency bands (N. D. R. Bhat et al. 2007). It is not yet clear whether nulling is an extreme form of mode-changing or an independent phenomenon (e.g. A. Naidu et al. 2017; A. Naidu et al. 2018).

Recent years have seen substantial progress in studies of pulsar nulling driven by the availability of sensitive single-pulse observations and increasingly sophisticated statistical methodologies. D. L. Kaplan et al. (2018) introduced a Gaussian mixture modelling framework for estimating nulling fractions directly from pulse-energy distributions, allowing null and burst states to be statistically separated without requiring hard thresholds. This method was also adapted and expanded by P. R. Brook et al. (2026) in a study of time evolution of nulling fractions in a sample of pulsars. G. Grover et al. (2026) analysed nulling pulsars using long-duration Murchison Widefield Array observations at 140–170 MHz, introducing the ‘ N_{sum} ’ method in which randomly selected groups of pulses are summed in order to statistically separate weak burst emission from the noise floor. The large number of pulses available per observation enabled detailed studies of nulling behaviour and quasiperiodic nulling phenomena.

Although there are many detailed single-pulse studies of bright individual pulsars, comparatively few studies have characterized pulse-to-pulse emission across large pulsar samples. Single-pulse studies of pulsars are inherently limited in integration time to a fraction of the pulse width, and therefore signal-to-noise ratio (S/N) is dominated by the instantaneous sensitivity of the telescope. This means that it is difficult to collect a sufficient quantity of high S/N data for a population of sources. However, the large collecting area and wide bandwidth of the MeerKAT radio telescope are particularly beneficial for single-pulse studies of pulsars. As a result, a key goal of the Thousand Pulsar Array (TPA) programme on MeerKAT was to produce a legacy dataset of ~ 1000 individual pulses from each of ~ 1000 pulsars (S. Johnston et al. 2020; X. Song et al. 2021).

The most recent large-scale surveys of pulse energy distributions were carried out with Murriyang, the Parkes radio telescope: that of 315 pulsars presented by S. Burke-Spolaor et al. (2012), observed as part of the Parkes High Time Resolution Universe survey, and that of 264 pulsars presented by M. B. Mickaliger et al. (2018) as part of the Parkes Multibeam Pulsar Survey. With six times the sensitivity and four times the number of pulsars,

the TPA single-pulse census represents a marked improvement in the overall number of pulsars for which high S/N measurements of individual pulsars have been made. Unlike survey-based data sets which typically have a fixed integration time, the TPA single-pulse census chose an integration time to capture at least 1024 pulses from most pulsars, with the exception of a small number of long-period pulsars (2 s), for which 512 pulses are recorded (see X. Song et al. 2023 for an in-depth description of the sample). This provides a broadly homogeneous dataset across a large fraction of the known pulsar population for studying short-term variability, complementing existing deep studies of individual pulsars with tens of thousands of pulses from a small number of pulsars (e.g. P. F. Wang et al. 2020), or vast numbers of pulses from millisecond pulsars (A. Parthasarathy et al. 2021; M. T. Miles et al. 2022).

In this paper, we present the MeerTime Single-pulse Pipeline (MTSP), developed specifically to process single-pulse observations from the MeerKAT telescope, and the resultant TPA single-pulse data set. The resulting calibrated and cleaned data for the full sample of pulsars are available via Zenodo,¹ and via a web interface² The single-pulse subpulse modulation properties of the data set have previously been published by X. Song et al. (2023), and the polarization properties of the pulsars with the brightest single pulses in the data set were studied by S. Johnston et al. (2024). Studies of drifting subpulses in individual pulsars have also been published from these data (A. Szary et al. 2022; J. A. Hsu et al. 2025; H. Wang et al. 2025). Here, we present an analysis of the pulse energy distributions of the TPA single-pulse data set. We use the large sample size to probe the evolution of nulling fraction and pulse energy distribution as a function of the pulsar’s period (P) and period derivative ($\dot{P}$).

In Sections 2 and 3, we explain how the data were recorded and processed with the MTSP. We present our methodology for calculating pulse energy distributions and nulling fraction in Section 4. Section 5 is devoted to presentation of the results of these measurements, discussion of them in the wider context of pulse single-pulse variability, and the conclusions are given in Section 6.

2 THE MEERTIME TIME-SERIES DATA RECORDING

The MTSP data are recorded as channelized time-series PSRFITS files (sometimes known as ‘filterbank’ or ‘search-mode’ data) using the Pulsar Timing User Supplied Equipment (PTUSE) hardware. A detailed description of the MeerTime observing system can be found in M. Bailes et al. (2020). For ‘L-band’ observations carried out as part of the TPA, the primary focus of this paper, the data are recorded with 0.8359375 MHz channels, sampled every $\sim 38.28 \mu\text{s}$. Observations for other MeerTime subprogrammes follow the same general principles, but may be recorded with higher or lower resolution, for example observations of millisecond pulsars are typically observed using four times as many 0.208984375 MHz channels sampled every $\sim 9.57 \mu\text{s}$.

The PTUSE computers split the observed data into two equal sub-bands, processed independently by one of the two GPUs in each machine. From 2019 May onwards the recording of the two sub-bands was synchronized so that the raw time-series data

¹<http://doi.org/10.5281/zenodo.18980771>

²<http://psrweb.jb.man.ac.uk/tpa/singlepulse>

Table 1. Time-series observing bands used by MeerTime. Data in 2019 March and April are recorded in two sub-bands, which are combined in subsequent processing. The number of frequency channels (N_{chan}) is given for the TPA project.

Dates	Centre freq (MHz)	Bandwidth (MHz)	N_{chan}
2019 Mar–2019 Apr	1283.58203125	856	1024
2019 May–2020 Feb 10	1283.58203125	642	768
2020 Feb 10–onward	1283.58203125	856	1024

could be combined during recording to produce a single PSRFITS file containing the entire observing band. However, for data recorded prior to 2019 May, the two sub-bands were not synchronized and so are combined during the single-pulse processing. Observations recorded in 2019 March may be affected by a buffer overflow error where the hardware was unable to keep up with the incoming data leading to either data packets being repeated or lost in the output data stream. This may result in very small spurious signals appearing, though the total packet loss is only around 4 parts in 10^9 . Nevertheless, to preserve data integrity, the recorded bandwidth was reduced to 642 MHz from 2019 April until 2020 February when we had confirmed that the issue was not present in the production version of the PTUSE hardware installed in 2019 December. The band configurations used at different times are listed in Table 1.

2.1 The Thousand-Pulsar-Array single-pulse data set

Single-pulse data are produced for essentially all TPA observations. In this paper, we focus on the set of observations determined by X. Song et al. (2023), which lists the ‘best’ single observation of each of 1192 pulsars. This is typically the longest observation, containing slightly more than 1024 consecutive pulses from the pulsar, and observed with close to the full MeerKAT array (i.e. no subarrays). Around 600 of these pulsars also have much shorter follow-up observations, unusually made with a subarray, with the aim to do timing and monitoring (see e.g. M. J. Keith et al. 2024). We also collect single-pulse data from these follow-up subarray observations, but due to disc space constraints we are unable to permanently archive the raw data and only the reduced data products are kept. In order to keep the scope manageable, in this paper we discuss only the analysis of the 1192 ‘best’ full-array observations.

3 THE MEERTIME SINGLE-PULSE PIPELINE

The MTSP was developed to process single-pulse data from the PTUSE backend, primarily the large volumes of TPA single-pulse data. The MTSP is written using PYTHON, and is designed to be run over a slurm cluster environment such as that on the OzStar supercomputer.³ The pipeline performs a number of operations:

- (i) Apply polarization calibration and other corrections needed for early science data.
- (ii) Detection and excision of narrow-band radio frequency interference (RFI).
- (iii) Identify on and off-pulse phases.

- (iv) Split the time-series into individual pulses at a suitable off-pulse phase.
- (v) Reduce the number of time and frequency samples through averaging.
- (vi) Generate diagnostic plots for quick visualization of the results.

In the default configuration, the output data have 1024 time samples per period, and 16 frequency channels. However, it is possible to adjust these settings to increase the time or frequency resolution if required for a specific science case.

3.1 Polarization calibration

For observations made with the L-band receiver recorded since 2020 April 11, polarization calibration is performed upstream before the data reach the PTUSE hardware, and so the data arrive at MTSP essentially fully calibrated. For data recorded before 2020 April 11, a polarization calibration is applied in the form of a per-channel Jones matrix supplied by the South African Radio Astronomy Observatory for each set of pulsar observations. Additionally, in order to form Stokes parameters, a correction for the receiver orientation is applied after calibration. Users of the raw time-series output of the PTUSE should note that the data may have an incorrect PSRFITS header parameter for the receiver orientation, specifically that `FD_HAND` should be set to `-1`. This is corrected automatically by the MTSP and it does not affect any downstream data products. See W. Straten et al. (2010) for a detailed explanation of the polarization conventions used by PSRFITS and PSRCHIVE.

3.2 RFI excision and on- and off-pulse windows

RFI is inherent to ground-based observations. The MTSP focuses exclusively on frequency-domain RFI mitigation by masking affected channels. We employ a fixed channel mask to target known, persistent interference, and an adaptive masking process based on off-pulse statistics. To ensure consistency for studies of pulse-to-pulse variability, the same frequency channels are removed from every pulse within a given observation.

The fixed mask (shown in the lower panel of Fig. 1) serves to increase consistency across the survey and prevents the statistics used in adaptive masking from being skewed by the most severely affected channels. These masked channels were selected based on the fraction of data masked in historical observations, and by association with known persistent RFI sources.

To adaptively identify RFI and define the on- and off-pulse regions, the time-series data are folded with DSPSR using the pulsar ephemeris and then fully averaged in time. A coarse first-stage cleaning is performed by fitting a noise-free template⁴ to each frequency channel using the standard Fourier convolution method implemented in PSRCHIVE (e.g. J. H. Taylor 1992). We define a test statistic for each channel based on the ratio of the resulting phase error to the S/N estimate; any channels where this statistic falls more than 1.5 times the interquartile range (IQR) from the median are excluded from the initial off-pulse determination.

The data are then frequency-averaged, and the template is re-fit to the profile to refine the on- and off-pulse regions. In the second cleaning stage, we calculate the off-pulse rms for every

³MeerTime data are housed and processed on the OzSTAR supercomputer at Swinburne University of Technology.

⁴See B. Posselt et al. (2021) and M. J. Keith et al. (2024) for a detailed description of template creation.




Figure 1. Fraction of data masked by the adaptive masking for each channel in observations during 2019 and 2020. Upper panel shows expected frequency bands for potential sources of interference. Lower panel shows the fixed channel mask applied to all data. Vertical dashed lines show the extent of data recorded in the 768-channel mode.

frequency channel (including those flagged in the first stage). A final mask is then applied to channels with an off-pulse rms exceeding 1.5 times the IQR. Following this final cleaning, the data are frequency-averaged once more to compute the definitive on- and off-pulse regions used for all later analysis.

3.3 Sub-band combination

The time-series data recorded by PTUSE during 2019 March and 2019 are recorded as two sub-bands split across two PSRFITS files, with a small but arbitrary start-time offset between the two band halves. In order to simplify later work, MTSP makes an effort to produce band-combined outputs that are largely identical to the outputs from later data. As the time-series are not exactly aligned in time, MTSP first forms de-dispersed (but not frequency-averaged) single-pulses from each sub-band, before combining them. In this way, the time bins are defined relative to the pulsar phase rather than the start of the observation, and so combination can be performed. The outputs from these combined observations can be identified as the band name will be set to either ‘1070+1498’ or ‘1123.5 + 1444.5’ in the MTSP metadata. Although all efforts are made to ensure these data are identical to those produced from later observations, we caution that there may be subtle differences in these data due to the different processing steps.

3.4 Creation of single-pulse data

In order to create the single-pulse data, the channelized time-series data are again processed by DSPSR in its single-pulse mode. The input ephemeris is modified to set the reference epoch (TZRMJD) such that DSPSR splits the pulses close to the centre of the largest off-pulse window. Once single pulses have been written, sub-band combination is performed (if required), polarization calibration is applied (if required), the data are corrected for dispersion and Faraday rotation, and converted to Stokes parameters. For TPA observations, the dispersion measure and rotation measure used are derived from the fold-mode TPA data (B. Posselt

et al. 2023). Finally, the pulses are averaged to 16 frequency channels, before being appended to an output PSRFITS file. Output data are split into files containing 256 consecutive pulses in order to reduce complexity of further data handling. A fully frequency-averaged version of the data are also created. The output data are produced as PSRFITS data files.

3.5 Data visualization

In addition to the PSRFITS data products, the pipeline also produces summary plots for each observation. These are intended as a quick-look tool, and do not represent the full breadth of analysis that is possible with the data. If required, the data are first averaged in phase such that the average S/N per on-pulse phase bin per period is at least 3. In order to improve the quality of the plots, the pipeline also tries to remove variations in the baseline over the observation, by subtracting a smooth model derived from a Gaussian process fit to the off-pulse data. These phase-averaged and baseline-subtracted data are also exported in Numpy (C. R. Harris et al. 2020) ‘.npz’ format in order to provide a quick and convenient data set without the need the specialist code needed interpret psrfits data. Summary plots over the whole data set are:

- (i) Histogram of phase-averaged intensity in the on- and off-pulse regions.
- (ii) Phase-resolved histogram of pulse intensity.
- (iii) Phase-resolved histogram of polarization position angle.
- (iv) Longitude-resolved fluctuation spectrum, i.e. the Fourier transform of each pulse phase bin.

And for each stack of 256 pulses, as a function of phase and pulse number:

- (i) Stokes parameters: I , Q , U , and V .
- (ii) Polarization position angle.
- (iii) Fractional linear and circular polarization.

Some example plots are shown in Fig. 2 for PSR J0034–0721, which is well known to have orthogonal polarization modes associated with its drifting subpulses (R. N. Manchester, J. H. Taylor & G. R. Huguenin 1975; C. D. Ilie et al. 2020).

4 MODELLING PULSE ENERGY DISTRIBUTIONS

Phase-averaged pulse energy distributions give a measure of the intensity distribution of individual pulses. This is useful for the broad understanding of the pulse-to-pulse behaviour of the pulsar, but also can be used to compute statistics that give an indication of the potential of the data set for further study of the pulsar behaviour.

4.1 Data preparation

To study the phase-averaged pulse energies, we take the frequency-averaged and baseline-subtracted output of the MTSP (the .npz files) and compute the pulse intensity time-series over the on-pulse regions. For comparison, we also take the average intensity in the off-pulse region of equal size. In the case of 36 pulsars where the on-pulse region is larger than the entire off-pulse region, the off-pulse statistics are estimated by the rms noise in the off-pulse time-averaged (i.e. folded) data products




Figure 2. Example summary plots for PSR J0034–0721 for a sequence of 256 pulses. Upper panel shows total intensity, and lower panel shows polarization position angle.




Figure 3. The phase-averaged intensity time-series for PSR J0034–0721 (thin blue line). The off-pulse values are shown for comparison (thicker orange line).

and scaled to estimate the expected frequency-averaged single-pulse rms. Time-domain RFI is excised by removing individual pulses where the rms of the entire off-pulse region falls outside 1.5 times the IQR. An example on-pulse time-series is shown in Fig. 3 for PSR J0034–0721, a pulsar which exhibits both nulls and very bright individual pulses, both of which can be seen in the figure.

4.2 Model fitting

In order to characterize the observed pulse energies, we fit a selection of models to the data. Four ‘simple’ single-component models are used, based on commonly observed pulse energy distributions (I. H. Cairns, S. Johnston & P. Das 2003; S. Burke-Spolaor et al. 2012):

(i) *Lognormal* – The intrinsic pulse energies are drawn from a lognormal distribution with parameters μ_L , the mean of the logarithm of the data, and σ_L^2 , the variance of the logarithm of the data.

(ii) *Normal* – The intrinsic pulse energies are drawn from a normal distribution with mean μ_N and variance σ_N^2 . In order to avoid significant negative pulses, the prior is restricted such that $\mu_N > 2.5\sigma_N$.

(iii) *Power law* – The intrinsic pulse energies are drawn from a power law with a probability density function given by

$$f(x) = (\alpha - 1)x_{\min}^{\alpha-1}x^{-\alpha}, \quad (1)$$

for $x \geq x_{\min}$, and $f(x) = 0$ otherwise, where α is the power-law index and $x_{\min}$ is the low-energy cut-off.

(iv) *Delta function* – The intrinsic pulse energies are of a constant value x_0 .

The delta-function model essentially sweeps up the cases where we have insufficient signal to determine properties of the underlying distribution.

Additionally, we fit for a number of ‘complex’ two- and three-component models:

- (v) *lognormal + lognormal*,
- (vi) *lognormal + normal*,
- (vii) *normal + normal*,
- (viii) *power law + lognormal*,
- (ix) *power law + normal*,
- (x) *lognormal + lognormal + lognormal*,
- (xi) *lognormal + lognormal + power law*,

where the + symbol indicates that pulses are drawn from one of the underlying distributions at random. These models allow the modelling to fit more complex pulse energy distributions, including those where there are multiple pulsar emission modes. The power-law components are required for cases where pulsars occasionally emit very bright pulses. All models also include a nulling fraction, which we define as the probability that a pulse is drawn from the off-pulse distribution instead of the on-pulse distribution.

The model pulse energy distribution, $f(E)$ is then given by the sum of the off-pulse distribution and the convolution of the off-pulse distribution with the intrinsic pulse energy distribution, weighted by the nulling fraction N_f . We can write this as

$$f(E) = N_f f_{\text{off}}(E) + (1 - N_f) \int f_{\text{int}}(\epsilon) f_{\text{off}}(E - \epsilon) d\epsilon, \quad (2)$$

where $f_{\text{int}}(E)$ is the pulse energy distribution of the non-null pulses [one of (i)–(xi) above] and $f_{\text{off}}(E)$ is the off-pulse distribution.

We solve for the model parameters, that is the parameters for each of the probability distributions and the nulling fraction, under a Bayesian framework. We compute the likelihood that the on-pulse data are drawn from $f(E)$, and the off-pulse data drawn from $f_{\text{off}}(E)$. That is, for N on and off-pulse measurements, the




Figure 4. Number of pulsars above a given cut-off in f_{13} , the fraction of the best-fitting pulse energy distribution above 3σ from the off-pulse mean. The dashed line marks the cut-off of $f_{13} > 0.3$ used in this paper.

log likelihood is given by

$$\ln \mathcal{L} = \sum_i^N \ln f(E_{\text{on},i}) + \sum_i^N \ln f_{\text{off}}(E_{\text{off},i}). \quad (3)$$

To explore the parameter space, we use the Dynesty nested sampler (S. Kposov et al. 2023). Unlike traditional Markov Chain Monte Carlo (MCMC) methods that map the posterior density directly, nested sampling is specifically designed for evidence calculation by systematically slicing the likelihood space into nested shells of known prior volume. This allows us to simultaneously estimate the posterior distribution of the parameters and reliably compute the Bayesian evidence, Z for the model. In the case where several models fit the data well, we select the model with the highest $\ln Z$ for the purpose of further study of the model parameters.

The overall approach is independently developed, but largely equivalent to the Gaussian mixture model presented by D. L. Kaplan et al. (2018) and the Bayesian nulling estimation method of P. R. Brook et al. (2026), and we would refer readers to those papers for discussion of the efficacy of the method. The primary difference in our model is that we select from a wide range of underlying pulse energy distributions rather than just a normal or lognormal distribution. A complementary approach to null-fraction estimation in low-S/N observations has recently been presented by G. Grover et al. (2026). Their method constructs distributions of randomly summed pulse energies, which statistically separate weak burst emission from the noise floor and allow null fractions to be inferred through Bayesian mixture modelling.

5 RESULTS AND DISCUSSION

5.1 Single-pulse detectability

Single-pulse S/N varies from pulse to pulse, therefore we use two metrics to quantify the fraction of detected single pulses, $f_{3\sigma}$, the fraction of pulses with phase-average intensity more than 3σ above that computed from the off-pulse, and f_{13} , the fraction of the best-fitting intrinsic on-pulse pulse energy distribution which is more than 3σ above the off-pulse mean. The values for $f_{3\sigma}$ will be lower as this includes nulls, whereas f_{13} only considers non-null pulses, but may be biased by inaccurate modelling of the underlying pulse energy distribution. In the sample, 367 pulsars have $f_{13} > 0.95$, i.e. 95 per cent of their pulses are intrinsically above 3σ , and 260 have $f_{3\sigma} > 0.95$. Fig. 4 shows the number of pulsars above any given f_{13} level in our sample.




Figure 5. Histogram of nulling fraction (N_f) for the sample. Solid histogram is for pulsars which show evidence for nulling, and unfilled histogram is for upper limits.

At low values of f_{13} the pulse energy distribution modelling becomes increasingly ambiguous and small differences in the on- and off-pulse noise levels can adversely affect the analysis leading to, for example, false detections of nulling.

To avoid the most ambiguous results, we applied a sample cut of $f_{13} > 0.3$. This value was chosen by inspection of individual low-S/N pulsars and via simulation. Simulations were performed by taking best-fitting pulse energy distributions of a small sample high-S/N pulsars and scaling them to lower flux density in order to investigate where the identification of nulling properties became unreliable. In principle, methods such as those in G. Grover et al. (2026) have been demonstrated to better quantify nulling in the low-S/N regime, and whilst we find it outside the scope of this paper we certainly encourage use of these methods in future work. In practice, we find that $f_{13} > 0.3$ produces reliable results, and therefore the remainder of the analysis will focus on the 749 pulsars for which $f_{13} > 0.3$.

5.2 Nulling fraction

Fig. 5 shows a histogram of the nulling fraction where detected at a $> 3\sigma$ level, as well as upper limits for all other pulsars. These results are also shown as a function of P and $\dot{P}$ on Fig. 6. There is a clear indication that the nulling fraction and number of nulling pulsars evolves across the P - $\dot{P}$ diagram. The nulling pulsars are mainly below $\dot{E} \sim 10^{33} \text{ erg s}^{-1}$ with a band of high null fractions appearing around $\dot{E} \sim 10^{31} \text{ erg s}^{-1}$. It is perhaps interesting that the ‘intermittent pulsars’ which exhibit very long-term nulls appear at the higher end of this $\dot{E}$ range ($\sim 10^{33} \text{ erg s}^{-1}$; A. G. Lyne et al. 2017), and whilst there are some high null-fraction pulsars around this region of the P - $\dot{P}$ diagram there does not generally seem to be any significant difference in the properties of our sample close to this $\dot{E}$.

In Fig. 7, we show the distribution of nulling fraction, as well as the fraction of pulsars showing significant nulling as a function of $\dot{E}$. Errors on the fraction of nulling pulsars are estimated by modelling the fraction as the success probability of a binomial distribution and deriving credible intervals from the corresponding beta posterior distribution (assuming a uniform prior). The general trend is a gradual decrease of nulling phenomena with increasing $\dot{E}$, though there is some hint that there is an excess of nulling pulsars at high $\dot{E}$. Inspection of the four nulling pulsars with $\dot{E} > 10^{35} \text{ erg s}^{-1}$ indicates they are all relatively low-S/N cases with only a statistical detection of nulling rather than clear periods of zero emission. Combined with the small numbers of




Figure 6. Nulling fraction as a function of P and $\dot{P}$ for the TPA single-pulse census. Lines of constant $\log[\dot{E}/(\text{erg s}^{-1})]$ and characteristic age are shown with dashed and dotted lines, respectively. Black dots are pulsars for which $f_{3\sigma} > 0.3$, and grey dots are for all other TPA pulsars. Known rotating radio transients (RRATs) and long-term nulling pulsars are overlaid with plus and cross symbols, respectively.




Figure 7. Nulling as a function of $\dot{E}$. Upper panel, the distribution of N_f for each $\dot{E}$ bin. Points indicate the median N_f . Lower panel, the fraction of pulsars exhibiting nulling. Circles show pulsars with $N_f > 0.01$ and squares show pulsars with $N_f > 0.15$. Error bars show 68 per cent bounds on the fraction of pulsars. Text values indicate the actual number of pulsars in each bin.

pulsars in these bins, we therefore suggest that this apparent rise at high $\dot{E}$ is not significant.

To investigate the evolution of nulling fraction on the P – $\dot{P}$ diagram in more detail, we can model it as being drawn from the Beta distribution, which is highly flexible for bounded random values.

We therefore model the N_f of any given pulsar as a random parameter drawn from a Beta distribution with shape parameter v , having probability density

$$f(N_f; \mu, v) = k N_f^{\mu-1} (1 - N_f)^{(1-\mu)v}, \quad (4)$$

where k is a normalization constant, and mean $\mu(P, \dot{P})$ given by,

$$\mu(P, \dot{P}) = m \left(\frac{P}{s} \right)^\alpha \left(\frac{\dot{P}}{10^{-15}} \right)^\beta. \quad (5)$$

The parameters α and β therefore represent the evolution of the mean nulling fraction over P and $\dot{P}$, respectively, and m as the typical N_f for a $P = 1$ s and $\dot{P} = 10^{-15}$ pulsar.

We then compute a log-likelihood summed over all pulsars p ,

$$\ln \mathcal{L}(\alpha, \beta, m, v) = \sum_p \ln [C_B(N_{u,p}, \mu_p, v) - C_B(N_{l,p}, \mu_p, v)], \quad (6)$$

where C_B is the cumulative distribution function of the Beta distribution, and $N_{u,p}$ and $N_{l,p}$ are 95 per cent upper and lower bounds on N_f for each pulsar. This is in effect reducing the N_f posterior estimates to a uniform distribution, and is done for computational simplicity. We then maximize the likelihood, solving for α , β , m , and v using common MCMC techniques provided by EMCEE (D. Foreman-Mackey et al. 2013). The results are $\alpha = 0.93 \pm 0.08$, $\beta = -0.10 \pm 0.03$, $m = 0.08 \pm 0.01$, and $v = 0.19 \pm 0.05$. This implies a nearly linear dependence of the nulling fraction with period, and little dependence on period derivative. The strong dependence of nulling fraction on spin period suggests that nulling is more closely linked to the large-scale structure of the magnetosphere than to available spin-down power. One possible interpretation is that nulling becomes more likely as the light-cylinder radius increases and the polar cap shrinks, potentially making the magnetosphere more susceptible to perturbations that intermittently disrupt radio emission. In this context, G. A. E. Wright (1979) and later J. M. Cordes & R. M. Shannon (2008) proposed that episodic accretion of circumstellar debris could trigger changes in magnetospheric state, leading to nulls or intermittent behaviour. While that model does not necessarily predict a scaling of infall rate with light-cylinder radius, the effectiveness of such external perturbations may plausibly depend on the size and geometry of the magnetosphere, which evolve strongly with spin period.

However, this is certainly not a unique explanation of the observed nulling behaviour, and these results should be considered over a wide range of theoretical models for pulsar nulling. Further, our analysis does not consider the time-scale of the nulling, which also is seen to evolve with P and $\dot{P}$ (A. G. Lyne et al. 2017), as does the apparent quasi-periodicity of null sequences in some pulsars (G. Grover et al. 2026), and there is evidence for decade-time-scale evolution of nulling properties in some pulsars (P. R. Brook et al. 2026).

5.3 Pulse energy distributions

We estimate the log evidence, $\ln Z$, for each of the pulse energy distribution models, and reject any models with $\ln Z - \ln Z_{\text{best}} < -2.5$. For the majority of pulsars this Bayesian model selection indicates more than one model that is compatible with the observed pulse energy distribution. 371 (51 per cent) of the pulse energy distributions are able to be explained with a single-component model and 360 (49 per cent) require at least two components. Although the very low-S/N pulsars are biased towards simple models, the fraction of pulsars requiring more than two components remains ~ 50 per cent even if we consider only the pulsars with $f_{3\sigma} > 0.95$. Fig. 8 shows two examples pulse energy distribution fits for pulsars with $f_{3\sigma} > 0.95$. PSR J1941–2602 shows a lognormal distribution with effectively no nulls. PSR J1507–6640




Figure 8. Visualizations of the pulse energy distribution of two example pulsars. Curves show samples from the posterior distribution for the preferred pulse energy distribution model, and histograms show the observed on- and off-pulse intensities. Intensity is normalized such pulses with the mean flux density have unit intensity. The inset box shows the number of pulses (N_p), the nulling fraction (N_f), and the mean flux density ($\bar{S}$). Left panel shows PSR J1941–2602, well fitted by a single lognormal component without any nulling. Right panel shows PSR J1507–6640, which requires at least two model components (shown here two lognormal components) and a nulling fraction of ~ 1 per cent. Similar figures for all pulsars are available in the associated data release.

shows both nulling (null fraction ~ 1 per cent) as well as an on-pulse energy distribution that requires at least two components, and the highest evidence is for the three lognormal component model.

With such a large sample of pulsars and multiple models that can fit many pulsars, we can try to simplify the data set by classifying the pulsars based on which models are compatible with the observed pulse energy distribution. Pulsars that are compatible with exactly one of the single-component models can be classed as lognormal (278 pulsars), normal (34 pulsars), or power-law (10 pulsars). Of the remainder, 49 pulsars are compatible with more than one single-component model, e.g. we cannot distinguish between lognormal and normal, and 366 are only compatible with the more complex multicomponent models, typically with more than one viable model. The remaining 12 pulsars are most compatible with the delta-function model and hence we are unable to discern anything about the underlying pulse energy distribution.

Given the large number of ‘complex’ pulse energy distributions, we can also classify pulsars based on any distribution types that are required across all preferred models for a given pulsar. For example if all preferred models require a lognormal distribution we classify this pulsar as ‘needs lognormal’. It is possible for a pulsar to require more than one type of distribution and therefore appear in multiple ‘needs’ categories. Fig. 9 shows the distribution of the pulsars requiring lognormal components across the P – $\dot{P}$ diagram (the other two categories are shown in Appendix A). The density of lognormal pulse energy distributions appears to increase with increasing $\dot{E}$. To investigate this further, Fig. 10 shows the number of pulsars in each ‘needs’ category as a function of $\dot{E}$. Although the numbers of pulsars in some bins is small, there does seem to be a subtle evolution of the pulse energy distributions, with more prevalence of lognormal components at high $\dot{E}$. There may also be a preference for power-law components around $\dot{E} \sim 10^{32} \text{ erg s}^{-1}$, and there is other evidence for complex emission behaviour in this $\dot{E}$ range. Long-term intermittent pulsars seem to lie along this $\dot{E}$ (A. G. Lyne et al. 2017), and there is evidence that there is a change in the pulse profile properties around $\dot{E} \simeq 10^{32.5} \text{ erg s}^{-1}$ (T. Olszanski et al. 2022). It is likely




Figure 9. Distribution of pulsars requiring lognormal components in their pulse energy distribution as a function of P and $\dot{P}$, shown with blue stars. Black dots are all pulsars for which $f_{33} > 0.3$, and grey dots are for all other TPA pulsars. Lines of constant $\log[\dot{E}/(\text{erg s}^{-1})]$ and characteristic age are shown with dashed and dotted lines, respectively.

that a phase-resolved study is necessary to properly understand if there is a physical relationship here.

5.4 Single-pulse luminosity

The single-pulse energy distributions are initially computed relative to the mean on-pulse energy. Although these individual single-pulse observations are not flux-calibrated, the corresponding folded profiles (which make use of the same underlying astrophysical signal) are automatically calibrated by the MeerTime pipeline. We can therefore approximate the absolute single-pulse energy in flux units by scaling the relative energies by the time-averaged flux density derived from the fold pipeline (B. Posselt et al. 2023).




Figure 10. Pulse energy distribution shape as a function of $\dot{E}$. This shows distribution shapes that are required for all best-fitting models, hence pulsars can lie in multiple shape categories, or none, and hence the sum of categories is not expected to sum to the total number of pulsars. Error bars show 68 per cent intervals estimated as in Fig. 7. The number of pulsars in each bin is shown as a fraction next to each point.

We convert this single-pulse flux density to an estimate of the bolometric radio luminosity by assuming a power-law spectral energy distribution with a constant spectral index of -1.2 between 10 MHz and 100 GHz. This effective spectral index was selected because it minimizes the systematic divergence in integrated luminosity across the overlapping sample of 42 pulsars shared with A. V. Bilous et al. (2020), who utilized a comprehensive spectral energy distribution model measured over a wide range of frequencies. Following the beaming model of A. Karastergiou et al. (2024), and utilizing the empirical pulse widths determined for the TPA sample (B. Posselt et al. 2021), the estimated bolometric radio luminosity is given by:

$$L = 4\pi D^2 S_0 \frac{\sin^2(\pi \delta_{10})}{\delta_{10}} \int_{10 \text{ MHz}}^{100 \text{ GHz}} \left(\frac{\nu}{\nu_0}\right)^{-1.2} d\nu, \quad (7)$$

where D is the distance to the pulsar, δ_{10} is the pulse width at 10 per cent of the peak intensity expressed as a fraction of the pulse period, and S_0 is the mean flux density measured at the reference frequency $\nu_0 = 1.28$ GHz. While these values provide a useful proxy for the true radio luminosity, we caution that significant unmodelled uncertainties remain. In particular, this metric relies on a highly simplified model of the spectral energy distribution, and the large distance uncertainties inherent to many pulsars in the sample will introduce corresponding scatter into the calculated luminosities.

Fig. 11 shows L as a function of $\dot{E}$ for the median flux density (i.e. that taken from the time-averaged observation) and from the brightest observed single pulse. We can see that many pulsars, especially those below $\dot{E} = 10^{32} \text{ erg s}^{-1}$ show individual pulses with luminosity close to or exceeding the long-term spin-down energy loss rate by 1–2 orders of magnitude, even where their time-averaged luminosity does not.

This suggests that the high luminosities ($\gtrsim \dot{E}$) observed in ultra-long-period transients such as GPM J1839–10 (N. Hurley-Walker et al. 2023) may not be unique to such objects. Indeed, extrapolating from the TPA population, we might expect a pulsar




Figure 11. Estimated radio luminosity of the TPA sample for the average flux density (plus symbol) and for the brightest observed pulse (stars), as a function of the inferred rotational energy loss rate $\dot{E}$. Dashed line shows the one-to-one correspondence of luminosity and $\dot{E}$.

with P and $\dot{P}$ similar to GPM J1839–10 would regularly show bright pulses with L_p in excess of $\dot{E}$ by a factor of 10 or more.

6 CONCLUSIONS

In this work, we have presented the TPA single-pulse data set and the MTSP pipeline used to produce this data set. The public data release of the ‘census’ observations is available at doi.org/10.5281/zenodo.18980771, and can also be accessed via a web interface at psrweb.jb.man.ac.uk/tpa/singlepulse which will be updated with future data releases in time. We calculate nulling fractions and model the single-pulse energy, which are also available online at psrweb.jb.man.ac.uk/tpa/pulse_energy and archived on Zenodo at doi.org/10.5281/zenodo.18982781.

The scale of the TPA single-pulse data set enables studies of the population-level statistics of single-pulse energy distributions at a new level of sensitivity. The results show that single-pulse energies evolve across the pulsar population, with some evidence of evolution of both nulling fraction and energy distribution shape with P or $\dot{E}$. However, they also show that the considerable complexity of pulsar radio emission variability is not sufficiently well captured by phase-averaged measurements. For example, measurements of phase-integrated pulse energy distributions cannot capture the fact that pulse profiles are made up of individual subpulses, which have complex behaviour such as subpulse drifting and mode-changing. Future work should therefore model phase-resolved single-pulse energies and investigate their relationships both to spin-down and to the diverse ways in which time-variability manifests in pulsar observations.

ACKNOWLEDGEMENTS

The MeerKAT telescope is operated by the South African Radio Astronomy Observatory (SARAO), which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. SARAO acknowledges the ongoing advice and calibration of global navigation satellite systems by the National Metrology Institute of South Africa (NMISA) and the time-space reference systems department of the Paris Observatory. Pulsar research at Jodrell Bank is supported by a con-

solidated grant (ST/X001229/1) from the UK Science and Technology Facilities Council (STFC). PTUSE was developed with support from the Australian SKA Office and Swinburne University of Technology. This work made use of the OzSTAR national high performance computing facility at Swinburne University of Technology. MeerTime data are housed on the OzSTAR supercomputer. The OzSTAR programme receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government.

DATA AVAILABILITY

The single-pulse data sets for the TPA single-pulse census used in this paper are archived via Zenodo, doi.org/10.5281/zenodo.18980771, and can also be accessed via a web interface at psrw.eb.jb.man.ac.uk/tpa/singlepulse. The results of the pulse energy distribution fitting are available at psrweb.jb.man.ac.uk/tpa/puls_e_energy archived on Zenodo, doi.org/10.5281/zenodo.18982781. The software used to generate the data is available at github.com/SixByNine/meertimesinglepulse

REFERENCES

- Backer D. C., 1970, *Nature*, 228, 42
- Bailes M. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e028
- Basu R., Mitra D., Melikidze G. I., 2017, *ApJ*, 846, 109
- Bhat N. D. R., Gupta Y., Kramer M., Karastergiou A., Lyne A. G., Johnston S., 2007, *A&A*, 462, 257
- Bilous A. V. et al., 2020, *A&A*, 635, A75
- Brook P. R., Gibson J. B., McLaughlin M. A., Surnis M. P., 2026, *MNRAS*, 547, stag260
- Burke-Spolaor S. et al., 2012, *MNRAS*, 423, 1351
- Cairns I. H., Johnston S., Das P., 2001, *ApJ*, 563, L65
- Cairns I. H., Johnston S., Das P., 2003, *MNRAS*, 343, 512
- Cairns I. H., Johnston S., Das P., 2004, *MNRAS*, 353, 270
- Cordes J. M., Shannon R. M., 2008, *ApJ*, 682, 1152
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Gajjar V., Joshi B. C., Kramer M., Karuppusamy R., Smits R., 2014, *ApJ*, 797, 18
- Grover G., Bhat N. D. R., McSweeney S. J., Lee C. P., Tan C. M., Fu S. C., Meyers B. W., 2026, *ApJ*, 997, 164
- Harris C. R. et al., 2020, *Nature*, 585, 357
- Hsu J. A., Weltevrede P., Wright G., Keith M. J., Oswald L. S., Song X., Wang H., 2025, *MNRAS*, 538, 1063
- Hurley-Walker N. et al., 2023, *Nature*, 619, 487
- Ilie C. D., Weltevrede P., Johnston S., Chen T., 2020, *MNRAS*, 491, 3385
- Jennings R. J. et al., 2024, *ApJ*, 964, 179
- Johnston S., Romani R. W., 2002, *MNRAS*, 332, 109
- Johnston S. et al., 2020, *MNRAS*, 493, 3608
- Johnston S., Mitra D., Keith M. J., Oswald L. S., Karastergiou A., 2024, *MNRAS*, 530, 4839
- Kaplan D. L., Swiggum J. K., Fichtenbauer T. D. J., Vallisneri M., 2018, *ApJ*, 855, 14
- Karastergiou A., Johnston S., Posselt B., Oswald L. S., Kramer M., Weltevrede P., 2024, *MNRAS*, 532, 3558
- Keith M. J. et al., 2024, *MNRAS*, 530, 1581
- Kikunaga T. et al., 2024, *Publ. Astron. Soc. Aust.*, 41, e036
- Koposov S. et al., 2023, *dynesty*. Zenodo
- Kramer M., Lyne A. G., O'Brien J. T., Jordan C. A., Lorimer D. R., 2006, *Science*, 312, 549
- Lundgren S. C., Cordes J. M., Ulmer M., Matz S. M., Lomatch S., Foster R. S., Hankins T., 1995, *ApJ*, 453, 433
- Lyne A. G. et al., 2017, *ApJ*, 834, 72
- Lyubarsky Y., 2019, *MNRAS*, 483, 1731
- Manchester R. N., Taylor J. H., Huguenin G. R., 1975, *ApJ*, 196, 83
- Mickaliger M. B., McEwen A. E., McLaughlin M. A., Lorimer D. R., 2018, *MNRAS*, 479, 5413
- Miles M. T., Shannon R. M., Bailes M., Reardon D. J., Buchner S., Middleton H., Spiewak R., 2022, *MNRAS*, 510, 5908
- Naidu A., Joshi B. C., Manoharan P. K., Krishnakumar M. A., 2017, *MNRAS*, 9, 1
- Naidu A., Joshi B. C., Manoharan P. K., Krishnakumar M. A., 2018, *MNRAS*, 475, 2375
- Olszanski T., Rankin J., Venkataraman A., Wahl H., 2022, *MNRAS*, 517, 1189
- Parthasarathy A. et al., 2021, *MNRAS*, 502, 407
- Philippov A., Uzdensky D. A., Spitkovsky A., Cerutti B., 2019, *ApJ*, 876, L6
- Posselt B. et al., 2021, *MNRAS*, 508, 4249
- Posselt B. et al., 2023, *MNRAS*, 520, 4582
- Smith F. G., 1973, *MNRAS*, 161, 9P
- Song X. et al., 2021, *MNRAS*, 505, 4456
- Song X. et al., 2023, *MNRAS*, 520, 4562
- Szary A., van Leeuwen J., Wright G., Weltevrede P., Agar C. H., Tiburzi C., Maan Y., Keith M. J., 2022, *ApJ*, 934, 23
- Taylor J. H., 1992, *Phil. Trans. R. Soc. Lond. A*, 341, 117
- van Straten W., Manchester R. N., Johnston S., Reynolds J. E., 2010, *Publ. Astron. Soc. Aust.*, 27, 104
- Wang H., Keith M. J., Weltevrede P., Wright G., Bhattacharyya B., Hsu J. A., Song X., 2025, *MNRAS*, 544, 234
- Wang P. F. et al., 2020, *A&A*, 644, A73
- Weltevrede P., Edwards R. T., Stappers B. W., 2006, *A&A*, 445, 243
- Weltevrede P., Stappers B. W., Edwards R. T., 2007, *A&A*, 469, 607
- Wen Z. G. et al., 2020, *ApJ*, 904, 72
- Wright G. A. E., 1979, *Nature*, 280, 40

APPENDIX A: ADDITIONAL $P-\dot{P}$ DIAGRAMS

Here, we show the distribution of the pulsars which require normal (Fig. A1) and power-law components (Fig. A2) across the $P-\dot{P}$ diagram.




Figure A1. As Fig. 9, but for pulsars requiring power-law components.




Figure A2. As Fig. 9, but for pulsars requiring power-law components.

This paper has been typeset from a $\text{T}_{\text{E}}\text{X}/\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$ file prepared by the author.