



HAL
open science

Simulation-based inference of radio millisecond pulsars in globular clusters

Joanna Berteaud, Christopher Eckner, Francesca Calore, Maïca Clavel, Daryl Haggard

► **To cite this version:**

Joanna Berteaud, Christopher Eckner, Francesca Calore, Maïca Clavel, Daryl Haggard. Simulation-based inference of radio millisecond pulsars in globular clusters. *Astrophys.J.*, 2024, 974 (1), pp.144. 10.3847/1538-4357/ad6b1e . hal-04600967

HAL Id: hal-04600967

<https://hal.science/hal-04600967v1>

Submitted on 1 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Simulation-based Inference of Radio Millisecond Pulsars in Globular Clusters

Joanna Berteaud^{1,2} , Christopher Eckner^{3,4,5} , Francesca Calore⁴ , Maïca Clavel⁶ , and Daryl Haggard^{7,8} ¹Department of Astronomy, University of Maryland, College Park, MD 20742, USA²NASA Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA³Center for Astrophysics and Cosmology, University of Nova Gorica, Vipavska 11c, 5270 Ajdovščina, Slovenia⁴LAPTh, CNRS, F-74000 Annecy, France⁵LAPP, CNRS, F-74000 Annecy, France⁶Univ. Grenoble Alpes, CNRS, IPAG, 38000 Grenoble, France⁷Department of Physics, McGill University, 3600 University Street, Montreal, QC H3A 2A7, Canada⁸Trottier Space Institute at McGill, 3550 University Street, Montreal, QC H3A 2T8, Canada

Received 2024 June 14; revised 2024 August 2; accepted 2024 August 3; published 2024 October 9

Abstract

Millisecond pulsars (MSPs) are abundant in globular clusters (GCs), which offer favorable environments for their creation. While the advent of recent, powerful facilities led to a rapid increase in MSP discoveries in GCs through pulsation searches, detection biases persist. In this work, we investigate the ability of current and future detections in GCs to constrain the parameters of the MSP population in GCs through a careful study of their luminosity function. Parameters of interest are the number of MSPs hosted by a GC, as well as the mean and the width of their luminosity function, which are typically affected by large uncertainties. While, as we show, likelihood-based studies can lead to ill-behaved posteriors on the size of the MSP population, we introduce a novel, likelihood-free analysis, based on marginal neural ratio estimation, which consistently produces well-behaved posteriors. We focus on the GC Terzan 5 (or Ter 5), which currently counts 48 detected MSPs. We find that 158^{+294}_{-104} MSPs should be hosted in this GC, but the uncertainty on this number remains large. We explore the performance of our new method on simulated Terzan 5-like data sets mimicking possible future observational outcomes. We find that significant improvement on the posteriors can be obtained by adding a reliable measurement of the diffuse radio emission of the GC to the analysis or by improving the detection threshold of current radio pulsation surveys by at least a factor of 2.

Unified Astronomy Thesaurus concepts: Millisecond pulsars (1062); Radio pulsars (1353); Bayes' Theorem (1924); Computational methods (1965); GPU computing (1969)

1. Introduction

Globular clusters (GCs) contain about half of the known population of Galactic millisecond pulsars (MSPs; Smith et al. 2023). Recently, the number of MSPs detected in GCs has risen rapidly thanks to new and powerful radio facilities like MeerKAT (Ridolfi et al. 2021, 2022) and FAST (Pan et al. 2021). Undoubtedly, MSPs are efficiently produced in GCs, which inform us about their possible formation channels. In addition, GCs have peculiar properties, namely a high stellar density and a profusion of old stars, which could enhance MSP formation (Ye et al. 2019). Some formation scenarios suggest that MSPs are old pulsars spun up through accretion in binary systems, while others favor former white dwarfs having undergone accretion-induced collapse. These different formation channels for MSPs also impact their abundances. However, to date, the number of MSPs in GCs remains highly uncertain. An alternative way to tackle the problem of their abundance is to study the luminosity function of MSPs in GCs, which is also impacted by their formation history.

The observation of MSPs in GCs, as well as elsewhere in the Galaxy, suffers from several biases. Brighter MSPs are easier to detect, and the period, dispersion, scattering or scintillation of their pulsed emission also impact their detection (Dai et al. 2017). As a result, data only reveal the most luminous sources,

likely in an incomplete fashion. Nonetheless, as the performance of radio instruments and their associated data sets improve, it will become easier to accurately extrapolate the luminosity function below the sensitivity of the observations.

Several works have investigated the luminosity function of MSPs in GCs in the past. Bagchi et al. (2011) used data from 85 MSPs in 10 GCs and, assuming a common luminosity function for all GCs, found that a log-normal distribution provides statistically better agreement with the data than the other distributions they tested. Their analysis relied on a comparison of Monte Carlo simulated MSP luminosities with observed ones by performing Kolmogorov–Smirnov and χ^2 tests. The analysis of Bagchi et al. (2011) relies on approximate Bayesian computation (ABC; Sisson et al. 2018), although the authors do not refer explicitly to this methodology. ABC tackles inference problems without a likelihood and estimates parameters by measuring a defined distance between simulated and real data. ABC belongs to the larger category of simulation-based inference (SBI; Cranmer et al. 2020), which has also been dubbed *likelihood-free* inference, as it does not rely on an explicit likelihood but on a data simulator instead. More recently, Chennamangalam et al. (2013) presented a *likelihood-based* Bayesian inference of the parameters of the MSP luminosity function, assuming again a log-normal distribution. Unlike Bagchi et al. (2011), each GC was treated independently from the others, ultimately producing different luminosity functions.

Since the publication of these studies, new MSP detections have been reported, and over the last decade inference

techniques have evolved. In particular, ongoing efforts at the intersection of machine learning and SBI have produced new tools (Cranmer et al. 2020) that have not yet been applied to the problem tackled by Bagchi et al. (2011) and Chennamangalam et al. (2013). The goal of this paper is to develop an up-to-date method to robustly determine the number of MSPs in GCs, along with its uncertainty, from individual detections using recently acquired data sets, and a new analysis framework: marginal neural ratio estimation (MNRE). Traditional inference methods, like the one mentioned above, produce a full, multidimensional, joint posterior, which is often not what one actually wants to study. Instead, marginalized and two-dimensional joint posteriors for specific parameters of interest are usually more interesting. Therefore, a large amount of computational time may be spent on solving the full problem while the scientific interest rather lies on a small subset of parameters. Marginal inference estimates the marginal posteriors of interest directly. The goals of the present work are to investigate the performance of SBI in comparison with traditional likelihood-based Bayesian analysis techniques applied to the same physical problem. Our target is to infer from synthetic radio MSP populations what quantities are essential to robustly constrain the population parameters in GCs and what properties help to obtain narrower, and yet well-behaved, posteriors of the inferred parameters.

In Section 2, we present the data sets, real and simulated, used in our analysis. Section 3 is dedicated to a discussion of the Bayesian analysis of Chennamangalam et al. (2013). In Sections 4 and 5, we present our SBI framework and its results. Discussion and conclusions are presented in Sections 6 and 7, respectively.

2. Data Set

2.1. Real Data Set

The basis for a statistical analysis of the intrinsic MSP population luminosity function is the measurement of the flux density, S , of a large sample of MSPs in GCs.

2.1.1. Pulsar Detection

In general, detection of pulsars relies on the measurement of a pulsed emission, not always but usually in the radio domain. A well-identified pulsation period is the key point of a detection, while a flux measurement is not essential and requires calibration beforehand. Indeed, when radio telescopes are used for pulsation searches, they record the *relative* intensity of the sky in the entire field of view of the telescope at a high sampling frequency. As a result, pulsars may lack absolute flux measurements despite their detection, and absolute flux measurements can vary from one instrument to another (see, e.g., the factor ~ 2 difference in the fluxes of Terzan 5 pulsars quoted by Bagchi et al. 2011 and Chennamangalam et al. 2013). Finally, as pulsation searches only record one-pixel images of the sky, the exact position of a pulsar cannot be deduced from a single observation, but requires tracking over several months. Hence, it is not always possible to associate a pulsar with a source seen in a radio image.

To be detected, the flux of a pulsar must lie above a certain flux threshold, which, in general, depends on the conditions along the line of sight as well as the radio telescope used for the observation. This threshold can be quantitatively stated via the radiometer equation (Dewey et al. 1985), which also depends

on the pulsation period P of the pulsar:

$$S_{\text{th}}(P) \propto \sqrt{\frac{w_{\text{obs}}}{P - w_{\text{obs}}}}, \quad (1)$$

where w_{obs} is the observed width of the pulse. A generic feature of the radiometer equation is that it predicts no universal detection threshold but a period dependence which flattens out toward long periods. Another parameter that strongly impacts the detection threshold is the column density of free electrons, ordinarily referred to as the dispersion measure, which increases the value of w_{obs} . For these reasons, an ideal pulsar data set for statistical analyses should ideally come from observations made with a single instrument in a unique configuration, for which both new discoveries and redetections (of known objects) should be reported. Despite ongoing efforts to collect data toward GCs through uniform, large surveys (see, e.g., Ridolfi et al. 2021, a census made with MeerKAT), to the best of our knowledge, a publicly available extended list of (re) detected pulsars in GCs quoting fluxes and telescopes' observing parameters does not exist. In what follows, we will therefore focus on a single GC, Terzan 5, for which recent, uniform flux measurements have been released.

2.1.2. Terzan 5

Terzan 5, or Ter 5, is the GC with the largest number of identified pulsars, including 49 with detected radio pulsations.⁹ Forty-eight of these have a period $P \leq 30$ ms and are therefore considered MSPs in the context of our work. The remaining pulsar, Terzan 5 J, has a period $P \simeq 80$ ms. The decade-old analysis of Chennamangalam et al. (2013) included 25 Terzan 5 pulsars (24 MSPs + Terzan 5 J) all with flux measurements. As the author noted, 34 pulsars were known at the time, but sources without flux measurements were excluded from the analysis.

Martsen et al. (2022) provided the most up-to-date flux and spectral index measurements of pulsars in Terzan 5. They identified 32 sources in archival data of the Green Bank Telescope (GBT), at 1.4 and 2 GHz, and computed their spectral indices, assuming power-law spectra. Ten additional fluxes were recorded with MeerKAT at 1284 MHz (Padmanabh et al. 2024). Thanks to the spectral index measurements, one can rescale the fluxes of Martsen et al. (2022) at the MeerKAT observing frequency. The final data set thus contains 42 pulsars with flux measurements, and seven pulsars without. This data set is *incomplete*, as not all of the sources detected in Terzan 5 have a flux measurement. When applying our MNRE framework to Ter 5, we will treat the Ter 5 data set without Terzan 5 J. Appendix A provides the flux densities used throughout this study.

Martsen et al. (2022) concluded that a dozen to more than a hundred additional pulsars are still to be discovered in Terzan 5. These subthreshold sources should contribute, at least partially, to the residual diffuse radio emission of Terzan 5, i.e., the total radio emission of the cluster minus the radio fluxes from the resolved point-like sources (and any other source of radio emission in the GC). Except for early observational results (Fruchter & Goss 2000)—most likely limited by the resolution of the instruments at that time—there are no recent, comparable assessments of the overall diffuse radio emission

⁹ <https://www3.mpifr-bonn.mpg.de/staff/pfreire/GCpsr.html>

from known Galactic GCs. Yet, this quantity, as we will see in what follows, can constrain the cumulative emission from subthreshold MSPs, and thus may help to narrow down the posterior distribution of the total number of sources in a GC.

2.2. Mock Data Set

Our ultimate goal is to infer the luminosity function of MSPs in GCs. To this end, we assume that the \log_{10} of the (pseudo-) luminosity L of MSPs in GCs follows a normal distribution, as suggested by previous studies (Bagchi et al. 2011; Chennamangalam et al. 2013):

$$f(\log_{10}(L)) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left[\frac{\log_{10}(L) - \mu}{\sigma}\right]^2\right). \quad (2)$$

L is expressed in millijanskys times square kiloparsec in Equation (2) and throughout this study. μ and σ are the mean and the width (i.e., the parameters to infer) of the luminosity function. Given the GC distance, fluxes are shifted according to

$$\log_{10}(S_i) = \log_{10}(L_i) - 2\log_{10}(d). \quad (3)$$

S_i and L_i are, respectively, the flux density and luminosity of pulsar i , and d is the distance to the cluster. d is expressed in kiloparsec and S_i in millijansky. We can define a probability density function $g(\log_{10}(S))$ that gives the probability of a pulsar in the cluster to have a flux S . Unlike f , this function is not unique and varies from one GC to another.

Considering N the *total* size of the MSP population in the cluster, we simulate the MSP population in the cluster by drawing N mock values of $\log_{10}(L)$ from Equation (2), and then we use Equation (3) to obtain, in turn, N mock values of $\log_{10}(S)$. Next, we identify the mock fluxes larger than a certain detection threshold value $S_{\text{th},i}$ and tag them as *detectable*. $S_{\text{th},i}$ may be different for each MSP (and is generally dependent on the cluster), and varies according to the radiometer equation. In the following, we consider two cases:

1. We assume that the detection threshold is cluster but not MSP dependent, as assumed by Bagchi et al. (2011) and Chennamangalam et al. (2013).
2. We mimic the effect of the radiometer equation by randomly drawing $S_{\text{th},i}$ from a half-normal distribution with mean value $S_{\text{th},\infty}$ representing the threshold for long-period pulsars. The width of the half-normal is taken to be $\sigma_{\text{th}} = S_{\text{th},\infty}$ such that larger detection thresholds are possible, accounting for possible systematical errors in the determination of the flux threshold from the radiometer equation, and mimicking an MSP-dependent threshold.

For a cluster, the number of *detectable* mock MSPs is $N_{\text{detected}} \leq N$. MSPs below the (case-by-case) threshold are counted as subthreshold sources and their cumulative flux S_{sub} is computed. Together with the flux of detectable MSPs, S_{det} , S_{sub} is responsible for the MSP radio emission of the cluster, S_{MSP} , which is at most equal to the total radio emission of the cluster, S_{tot} :

$$\begin{aligned} S_{\text{tot}} &\geq S_{\text{MSP}} = \sum_i S_i = S_{\text{det}} + S_{\text{sub}} \\ &= \sum_{S_i \geq S_{\text{th},i}} S_i + \sum_{S_i < S_{\text{th},i}} S_i. \end{aligned} \quad (4)$$

In ideal cases, the value of each $S_i \geq S_{\text{th},i}$ should be known. However, as mentioned in Section 2.1.1, this might not always

be the case, and may potentially hamper our ability to infer the posterior distributions of the MSP population. To resemble actual data, we introduce a parameter, p_{fluxless} , which represents the proportion of MSPs for which we have a detection but no flux measurement. $p_{\text{fluxless}} = 1 - m/N_{\text{detected}} \leq 1$, where m is the number of detected MSPs which have a flux measurement. In our simulation, we guarantee that the fraction of fluxless MSPs matches the one in real data. Finally, we define

$$S_{\text{diff}} = S_{\text{tot}} - \sum_m S_i, \quad (5)$$

which corresponds to the diffuse radio flux of the GC. S_{diff} may originate also from sources other than MSPs so that $S_{\text{diff}} - S_{\text{sub}} \geq 0$. The diffuse flux measurement is nothing but an upper limit on the total flux from the unresolved part of the GC's MSP population.

Our mock data set realizations are based on the Terzan 5 best-fit results of Chennamangalam et al. (2013), as we detail below.

3. Bayesian Analysis

3.1. Framework

Our Bayesian analysis relies on the framework described in Chennamangalam et al. (2013), which is briefly summarized here. Let M be our model, θ its parameters, and D the data. Then, according to Bayes' theorem:

$$P(\theta|D, M) \propto \mathcal{L}(D|\theta, M)p(\theta|M), \quad (6)$$

where $P(\theta|D, M)$ is the posterior distribution of the parameters given the data and the model, $\mathcal{L}(D|\theta, M)$ is the likelihood of the data given the parameters and the model, and $p(\theta|M)$ is the prior distribution of the parameters given the model. The proportionality sign accounts for the evidence $p(D)$ that divides the right-hand side of the equation, which, as a constant, can be ignored here. The likelihood $\mathcal{L}(D|\theta, M)$, associated with a given cluster, takes contributions from the following three independent likelihoods:

1. The likelihood of having a set of pulsars with fluxes $\{S_i\}$, computed as the product of the $g(\log_{10}(S_i))$.
2. The likelihood of observing N_{detected} (or m) pulsars in a cluster with N pulsars, computed as a binomial distribution.
3. The likelihood of observing a total flux S_{tot} , which has a Gaussian distribution around $N\langle S \rangle$, where $\langle S \rangle$ is the average MSP flux in the cluster. Here, as done in Chennamangalam et al. (2013), we assume $S_{\text{tot}} = S_{\text{MSP}}$.

Our luminosity model then has five parameters: $\theta = \{N, \mu, \sigma, S_{\text{th}}, d\}$. We adopt uniform priors for the first four parameters. A Gaussian prior is chosen for d , reflecting the idea that the distance to the GC is already known with some uncertainty. All priors are independent and similar to the ones used by Chennamangalam et al. (2013). Table 1 summarizes the prior sampling and ranges. Focusing on Terzan 5, the maximal number of MSPs in the cluster is set to 500, and the mean and the width of the distance prior are 5.5 and 0.9 kpc, respectively, according to the result of Ortolani et al. (2007).

Table 1
Summary of the Applied Prior Ranges for the Bayesian Analysis (Section 3) and the `swyft` Simulator (Section 4) of the GC MSP Population

Bayesian			swyft		
Parameter (1)	Sampling (2)	Prior Range (3)	Sampling (4)	Prior Range (5)	Physical Prior Range (6)
N	Uniform	$[N_{\text{detected}}, N_{\text{max}}]$	Log-uniform	$[\log_{10} N_{\text{detected}}, 2.7]$	$[N_{\text{detected}}, 500]$
μ	Uniform	$[-2.0, 0.5]$	Uniform	$[-2.0, 0.5]$	$[10, 3162] \mu\text{Jy}$
σ	Uniform	$[0.2, 1.4]$	Uniform	$[0.2, 1.4]$...
d	Normal
$S_{\text{th},\infty}$	Uniform	$[0, \min(S_i)]$	Log-uniform	$[0.5, 1.6]$	$[3, 40] \mu\text{Jy}$

Note. The priors used in the Bayesian analysis are the ones used by Chennamangalam et al. (2013). The sixth column states the prior ranges of the fifth column in physical units, while the fifth column is the numerical input for the respective probability distribution functions.

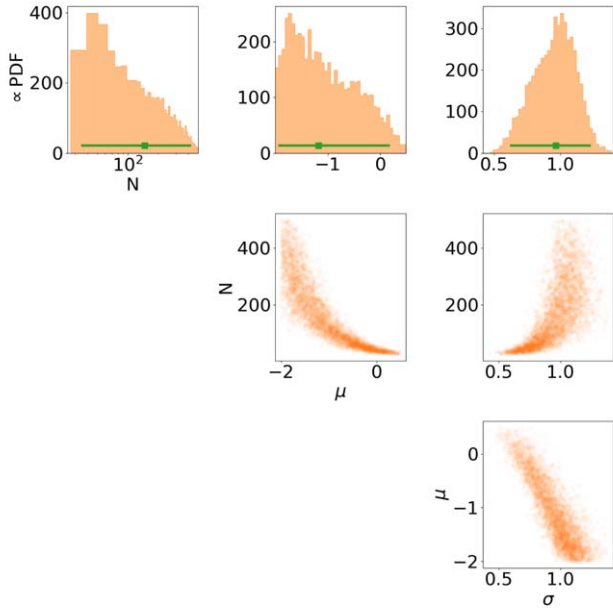


Figure 1. Reproduction of the likelihood-based analysis results of Chennamangalam et al. (2013) for Terzan 5. The top row shows the marginal probability density functions (PDFs) in orange and the median and its 95% error in green for each parameter. The position of the samples in each two-dimensional parameter space are also shown (second to last rows).

3.2. Results

We perform Bayesian analysis using `PyMultiNest`, a Bayesian inference tool written for `Python` (Buchner et al. 2014), which samples the parameter space using a Monte Carlo algorithm based on nested sampling. While this method samples parameter space more efficiently compared to other methods, the computation time still increases with the number of live points n_{lp} , i.e., the size of the set of samples drawn from each prior. To identify the number of live points needed for our analysis, we first reproduce the results of Chennamangalam et al. (2013). We strictly follow their analysis framework, summarized in Section 3.1, and apply it to the Terzan 5 pulsar sample; see Section 2.1.2. For a meaningful comparison, we used the fluxes of the 25 Terzan 5 MSPs in Table 1 of Chennamangalam et al. (2013). We conclude that a number of live points $n_{\text{lp}} \sim 1000$ suits our analysis. Our posteriors, shown in Figure 1, clearly illustrate that N , μ , and σ are degenerate to some extent, as already noted by Chennamangalam et al. (2013).

The 95% credible interval on N obtained through the likelihood-based analysis is not significantly narrower than the prior chosen for the parameter (e.g., 32–452 versus 25–500). This can indicate that the data are not informative enough to constrain the luminosity function and the size of the MSP population in Terzan 5.

This hypothesis can be tested by assessing the coverage of the credible intervals: In case of correct coverage, the $x\%$ credible interval of a parameter should contain the true parameter value in $x\%$ percent of the cases. If the $x\%$ credible region contains the true parameter value in $y\%$ percent of the cases, where $y > x$, the credible interval is too wide and it is said to be *conservative*. On the other hand, if $y < x$, the interval is too narrow and it is said to be *overconfident*. Our goal in this section is to assess the quality of the coverage obtained with the Bayesian analysis, which ultimately informs us about the correctness of the statistical inference. If the data are genuinely not informative enough to constrain the number of MSPs in Terzan 5, we should find $x \simeq y$ for all values of x . We simulate 500 mock data sets according to Section 2.2 using as true parameters of the model some of the results obtained by Chennamangalam et al. (2013) for Terzan 5. Namely, we assume $N = 142$, $\mu = -1.2$, and $\sigma = 1.0$. The sensitivity threshold is set to $S_{\text{th}} = 0.02$ mJy and the distance to $d = 5.5$ kpc. N_{detected} varies from one simulation to another. Finally, as in Chennamangalam et al. (2013), we have $p_{\text{fluxless}} = 0$, i.e., $m = N_{\text{detected}}$. We apply the Bayesian analysis streamlined in Section 3 to the mock data sets and use the posteriors of N to assess the coverage of this parameter. We build a symmetric interval around the median N_{med} of the posterior distribution $P(N)$ such that the integral of the posterior from the lower bound of the interval N_{low} to the median equals the integral of the posterior from the median to the upper bound of the interval N_{high} . We start from $N_{\text{low}} = N_{\text{high}} = N_{\text{med}}$ and increase the size of the interval until $N_{\text{low}} = 142$ or $N_{\text{high}} = 142$. We compute

$$\begin{aligned}
 x &= 100 \times \int_{N_{\text{low}}}^{N_{\text{high}}} P(N) dN \\
 &= 200 \times |C(142) - C(N_{\text{med}})|,
 \end{aligned} \tag{7}$$

where C is the cumulative distribution function associated with $P(N)$. When the true value is close to (far from) the median value, x takes a low (high) value. By definition, x cannot be larger than 100. If x_k is the value that x takes in simulation k , y takes the values y_k , computed as the percentage of values of y among all simulations that verify $y \leq y_k$. Hence, the y_k reflect the cumulative distribution of the x_k . The values taken by x and

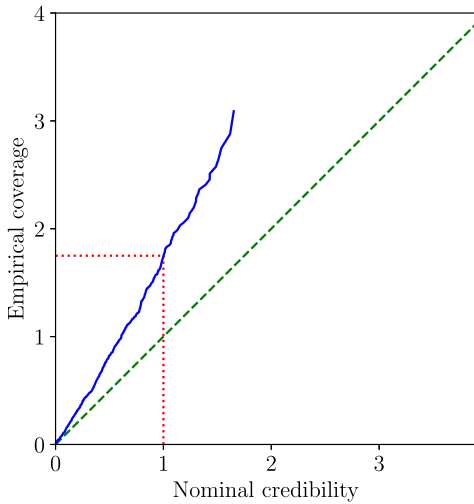


Figure 2. Coverage (blue line) of N obtained from the application of the Bayesian analysis to 500 mock data sets ($N = 142$, $\mu = -1.2$, and $\sigma = 1.0$). Ideally, the coverage line should follow the green dashed diagonal line. Here, the plot tells us that the method produces conservative results, i.e., that the posteriors are too wide. The red lines illustrate that the 68% (1σ) confidence interval contains the true value in more than 68% of the cases.

y are converted from percentage to significance level, which is the inverse of the error function of the percentage divided by a square root of 2. This conversion is convenient for plotting and interpreting the coverage. A perfectly calibrated posterior follows a diagonal in the space of NOMINAL versus EMPIRICALLY derived coverage. While this coverage test is indicative of ill-calibrated posteriors that necessitate a redesign of the inference pipeline, a positive outcome for this test does not imply that the obtained posterior distributions are optimal given the available statistical power of the data set.

Our results, shown in Figure 2, indicate that the posteriors of N tend to be too conservative. For all of them, the true value of $N = 142$ falls within the 92.7% credible interval (1.79 confidence level), while this should only happen for 92.7%, i.e., 463, of them. The coverage at higher confidence level cannot be estimated because the true value never falls within the 99% credible interval but outside of the 92.7% confidence interval. Therefore, a different analysis can obtain narrower posteriors with the actual data, while still being statistically correct. We note that the general aspect of the coverage may depend on the values chosen for the true parameters. What our results tell us is, should the actual parameters of Terzan 5 be the one we chose (see also Section 5), there is a high chance that the Bayesian analysis would yield posteriors that are too wide for the number of sources in the cluster.

4. Swyft Implementation on MNRE

Marginal (neural) ratio estimation focuses on the likelihood-to-evidence ratio, which, according to Bayes’ theorem, equals the posterior-to-prior ratio:

$$\begin{aligned} r(\theta, D, M) &= \frac{\mathcal{L}(D|\theta, M)}{p(D)} = \frac{p(\theta|D, M)}{p(\theta|M)} \\ &= \frac{p(D, \theta)}{p(D)p(\theta)}. \end{aligned} \quad (8)$$

Let us now define a binary variable Y , such that $Y = 1$ when pairs (D, θ) are jointly drawn from $p(D, \theta)$, and $Y = 0$ when

pairs are marginally drawn from $p(D)p(\theta)$. One can show that training a binary classifier f_ϕ for Y is equivalent to learning the ratio $r(\theta, D, M)$.

We rely on a specific implementation of MNRE: *swyft*. This is a simulation-based inference tool that uses MNRE methods, i.e., its classifier f_ϕ uses a neural network and learns from mock data produced by a simulator (Miller et al. 2021, 2022). *swyft* has been proven to be efficient at inferring cosmological parameters from cosmic microwave background measurements with posterior convergence achieved using orders of magnitude fewer calls than Markov Chain Monte Carlo methods (Cole et al. 2022). More recently, Anau Montel et al. (2023) showed that the estimation of the warm dark matter mass from strong lensing images could also benefit from MNRE. Our *swyft*-based code, described in Sections 4.1 and 4.2, is publicly available online (Berteaud et al. 2024).¹⁰

4.1. Simulator

In the framework of SBI, we are not bound to limit the number of model parameters and construct a likelihood function in the first place. In particular, the inclusion of nuisance parameters is facilitated. Hence, we extend the mock data generation outlined in Section 2.2 to profit from the capabilities of SBI. Our *swyft* radio MSP simulator generates mock data for a single GC. The baseline simulator depends on the parameters described in Section 2.2, whose priors are varied as reported in Table 1. We adopt a varying flux threshold, implemented as a nuisance parameter, which mimics an MSP-dependent threshold.

Simulator. This simulator yields a catalog of detected MSPs as well as the additional cumulative flux of the population too dim to be resolved individually. The information about the flux of each detected source is stored in an array, sorted in descending order. The length of the array is set to a fixed number (the upper bound of the prior on N). Array entries without a detected source are assigned the fill value “-1.” As we can have $p_{\text{fluxless}} \neq 0$, the flux value is replaced by the fill value “-1” for above-threshold, fluxless sources. We emphasize that all detected sources, including the fluxless ones, are assumed to be above the flux detection threshold, and vice versa. The simulator generates a second array that states N_{detected} subtracted by the cumulative sum of detected sources.

Prior choices. We aim to infer the posterior distributions for the model parameters based on the set of priors summarized in Table 1. These priors are slightly wider than what was used in Chennamangalam et al. (2013) and in Section 3 except for the prior on σ , which we adopt from this earlier work. Previous scans of the luminosity functions’ parameter ranges seem to support this prior range (Bagchi et al. 2011). We note that our objective is to infer information from simulated data so that we are rather agnostic about the frequency at which we sample the luminosity function and at which we define the detection threshold. Note that, in contrast to Section 3, we use a log-uniform prior on the number of MSPs N . We made this change to more evenly sample the larger range of up to 3000 sources instead of 500. When we later analyze the real sample of Terzan 5 (see Section 5.2), we will adopt the same prior definition.

¹⁰ [10.5281/zenodo.11403867](https://doi.org/10.5281/zenodo.11403867)

Table 2Summary of Network Architecture and `swyft` Hyperparameters Used in This Work, as Outlined in Section 4.2

Hyperparameter	Value (w/o Dif-fuse Flux)	Value (w Dif-fuse Flux)
ResNet: # input features	(500/500)	(500/500/500)
ResNet: # output features	(32/16)	(32/16/32)
ResNet: # hidden features	(128/128)	(128/128/128)
ResNet: # blocks	(2/4)	(2/4/5)
Sample size	10^5	
Training-to-validation ratio	70:30	
Training/validation batch size	64/64	
Optimizer	Adam	
Initial learning rate	8.5×10^{-4}	
Learning rate scheduler	ReduceLROnPlateau	
Learning ratio schedule decay factor	0.3	
Learning ratio schedule patience	5	
Maximal number of training epochs	100	
Early stopping patience	20	
Noise resampling	true	

Note. For the ResNet hyperparameters, the numbers in parentheses denote the values selected for each data column.

4.2. Architecture

As described in Section 4, the parameter inference in neural ratio estimation relies on training a classifier f_ϕ that allows for discriminating samples drawn from a joint or a marginal probability distribution. This classifier is a neural network in the case of `swyft`. Hence, design decisions for the network are heavily influencing the robustness, performance, and quality of the inference. We have to choose a network architecture that is suitable for the problem at hand and adequate for the data format generated by the simulator, which itself is motivated by the structure of real observations of radio MSPs. In our case, we obtain a catalog of detected sources represented by a two-dimensional array as detailed in Section 4.1; in abstract terms, a matrix with a certain number of rows and columns.

Inference network architecture. This data structure allows us to resort to network architectures typically employed in image processing. To increase the flexibility of the network, we choose the deep neural network architecture of the ResNet (He et al. 2015), which has been used in the somewhat similar context of source detection in gamma-ray data of the Fermi Large Area Telescope (Anau Montel & Weniger 2022; Horangic et al. 2023). In particular, we utilize the ResNet implementation natively provided in `swyft`. Each column of the data array is first passed through a normalizing layer and afterwards processed by a dedicated ResNet (with different structural properties) that takes one-dimensional vectors as input. The chosen ResNet hyperparameters appear in Table 2. In Figure 3, we provide a visualization of the ResNet’s core structure. The output vectors of each ResNet are concatenated to a single vector, which is subsequently passed to the default ratio estimator implementation of `swyft`. Thus, the initial data processing with ResNets compresses the data and generates a summary statistic learned and optimized during the network’s training.

Swyft hyperparameter settings. Besides the structural hyperparameters of the ResNet, training and inference with the MNRE algorithm of `swyft` requires the specification of further parameters. A summary of the parameter choices is provided in Table 2. A subset of the stated parameters controls the training procedure; the sample size refers to the total number of simulated radio MSP populations for the training. This sample is split into a subset used for training (70%) and a validation data set (30%) used to evaluate the performance of the trained network on data it has not seen during the training iterations. Both data sets are fed into the network in batches of 64 samples. The network training is performed with the Adam optimizer in, at most, 100 epochs starting with an initial learning rate of 8.5×10^{-4} . This learning rate may decrease during the training if after five consecutive epochs (learning ratio schedule patience) no improvement in the training loss was achieved. In this case, the current learning rate is reduced by 30%. If 20 consecutive epochs did not result in an improvement in the training loss (early stopping patience), the training terminates immediately. Lastly, we allow for noise resampling during the training process. The idea is to regenerate the noise in the data while keeping the training samples the same. In practice, the noise is represented by the set of detected sources without flux measurements. For each sample, we repeat per epoch the selection of MSPs that are listed with a flux value as the simulator has stored the full catalog of detected sources with their original randomly drawn fluxes. This procedure effectively enhances the number of training samples since the data set never looks like the one used in the previous epoch. It has been shown that noise resampling is a crucial way to prevent overfitting the training data and it stabilizes the shape of the posterior distributions (Alvey et al. 2023a, 2023b).

Including diffuse flux measurements. Motivated by the approach in Chennamangalam et al. (2013), we prepare an extension of the basic formalism outlined above. We aim to explore the impact of an additional diffuse radio flux measurement S_{diff} by extending the simulator: We create a third column of the generated catalog that utilizes the already computed cumulative flux of subthreshold sources S_{sub} . All rows with no entry or a subthreshold source are assigned the value $S_{\text{tot}} - S_{\text{sub}}$. Subsequently, following the ordering of flux in the first column, this value is decremented by the flux of the respective detected MSP. Sources without a flux measurement are included with zero flux.

Swyft hyperparameters with diffuse flux. This scenario does not require much tuning of the hyperparameters chosen in the baseline setup, but we must add a third ResNet that accepts the list of fluxes describing diffuse and resolved contributions. The selected parameters are stated in the third column of Table 2.

5. Swyft Results

Now that we have established the simulator of radio MSPs as well as the inference approach, we aim to study the performance of the method and the properties we are able to extract given a realization of synthetic target data. The idea is to choose the underlying model parameters to well represent what an observationally obtained radio MSP catalog looks like.

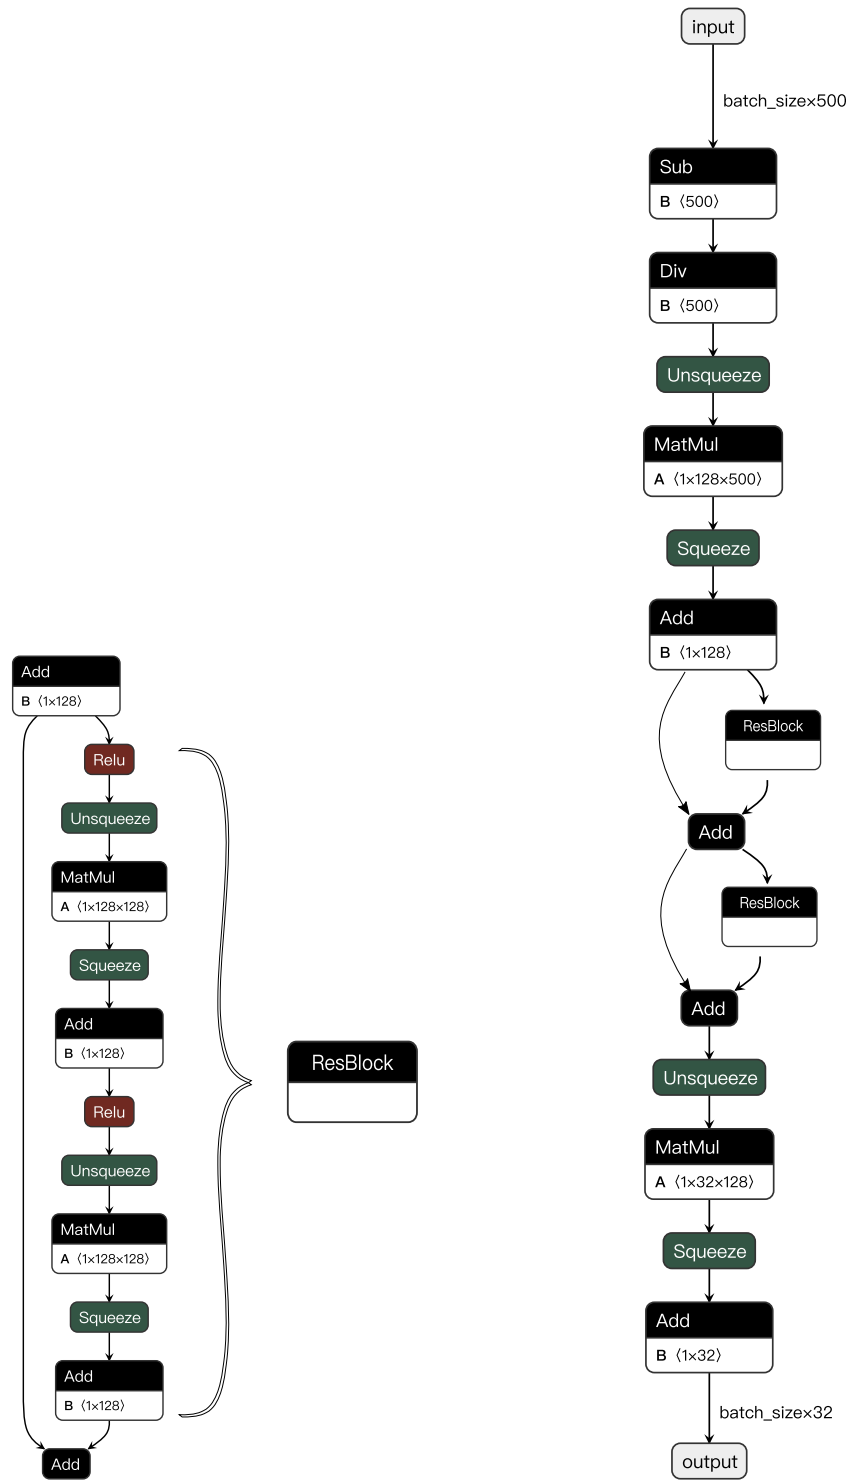


Figure 3. Flow chart of the ResNet implementation used in this work to generate summary statistics for the *swyft* log-ratio estimator. Left: the general structure of a block “ResBlock” of the ResNet used as the core ingredient to the overall ResNet layout. The operation Add performs an element-wise binary addition on a vector of a given size. MatMul refers to a matrix multiplication based on the dimensions specified in the box. Unsqueeze and squeeze augment or reduce the dimension of the input vector. Relu is the standard activation function known as a rectified linear function, $y = \max(0, x)$, applied element-wise to an input vector. Right: full structure of a ResNet defined with two “ResBlocks” following the hyperparameters listed in Table 2 for the case without S_{diff} summarizing the information in the list of detected MSP fluxes. The operations Sub and Div perform the normalization of the input flux vector via subtraction of the mean and division by the total flux. The other ResNet structures are similar in shape with altered numbers of “ResBlocks” and input and output features.

5.1. Mock Data

We examine two distinct mock MSP populations. First, we make the direct connection to the Bayesian analysis of Section 3.2 and adopt the same defining population parameters.

Then, we turn toward a mock data definition that is oriented closer to the current observational data for Ter 5. The latter case serves as a testing ground to explore the capabilities of our simulator and SBI. We investigate different scenarios for the

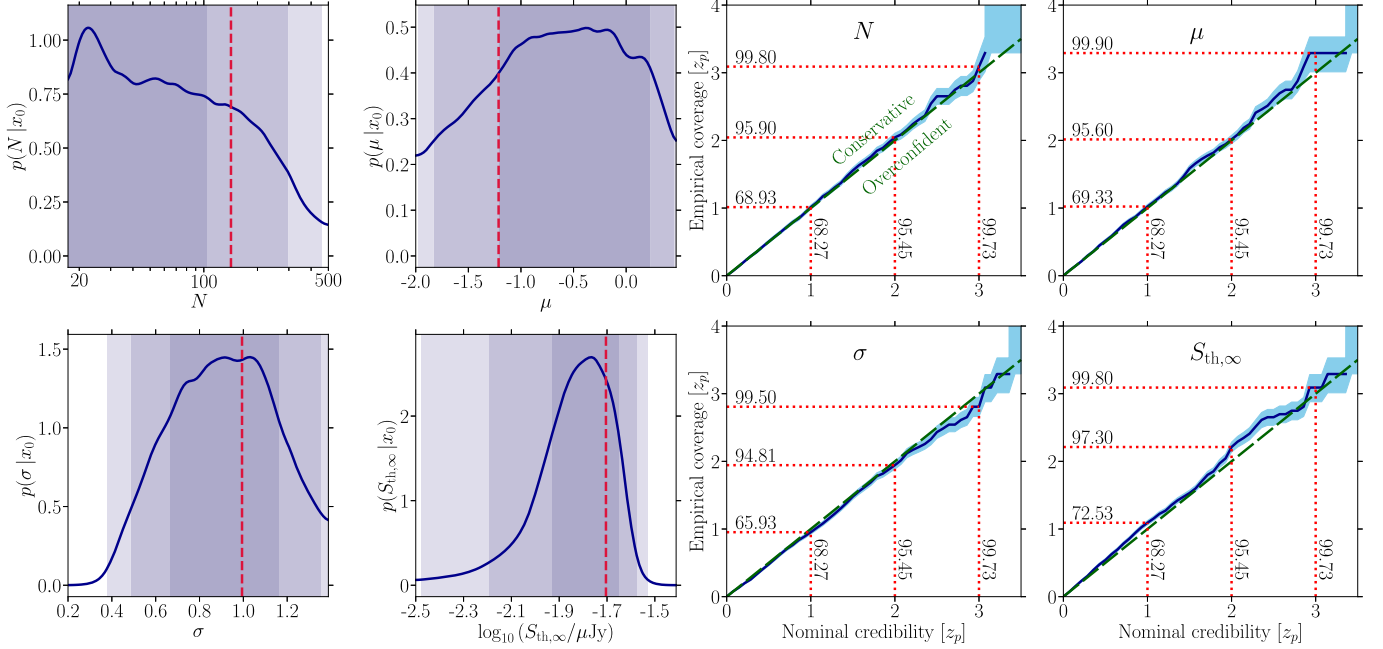


Figure 4. Left: complete set of one-dimensional marginal posterior distributions (blue curves) for the four parameters characterizing the radio MSP population model similar to the one found in Chennamangalam et al. (2013) with parameters $N_{\text{detected}} = 17$, $N = 142$, $\mu = -1.2$, $\sigma = 1.0$, $S_{\text{th},\infty} = 0.02$ mJy, and $p_{\text{fluxless}} = 0\%$. The corresponding 1σ , 2σ , and 3σ contours are displayed as shaded bands in decreasing opacity. The vertical dashed red lines denote the true parameters of the target observation. Right: *swyft* coverage results for the mock data set of the left panel. The horizontal axis states the nominal credibility interval (in significance) while the vertical axis shows the empirically determined coverage. The dashed green line indicates perfect coverage while the solid blue line is the obtained average coverage; the light blue band denotes the 68% containment band derived from 1000 simulations.

parameters impacting the quality of the inference, i.e., the percentage of detected sources without flux measurements, p_{fluxless} , the availability of diffuse measurements, and the detection threshold at large periods, $S_{\text{th},\infty}$. For what follows, the *swyft* analysis is entirely based on the simulator described in Section 4.1 with priors stated in Table 1 and the inference architecture of Section 4.2 adhering to the specifications in Table 2.

5.1.1. Mock Population Following Chennamangalam et al. (2013)

As a reminder, the MSP population in this case is characterized by $N = 142$, $\mu = -1.2$, $\sigma = 1.0$, $S_{\text{th},\infty} = 0.02$ mJy, and $p_{\text{fluxless}} = 0\%$. We generate 1000 realizations of this MSP population and determine the mean expected number of detected MSPs (while keeping the size of the population the same). We obtain $N_{\text{detected}} = 17$, which we fix as a hyperparameter of our simulator.

We visualize the results of the parameter inference with *swyft* in the left panel of Figure 4. The plot depicts the one-dimensional marginal posterior distributions (blue line) for all four model parameters with corresponding 1σ , 2σ , and 3σ contours highlighted as shaded bands with decreasing opacity. The injected parameter values are marked with red dashed lines. The training and subsequent parameter estimation lasted around 1.5 hr on a single NVIDIA A100 GPU. A direct comparison with Figure 1 reveals a striking similarity of the shape of the posterior distributions for N and μ and σ . Ultimately, both methods produce almost the same qualitative results, which is reassuring in the sense that the available information is equally well translated into posterior distributions. The median and the 95% credible region of N are reported in the first line of Table 3.

Table 3

Median and 95% Credible Interval of the Number of MSPs N Obtained through the *swyft* Analysis Applied to Mock Data for Various Cases

Case	Injected N	Inferred N
(1) $N_{\text{detected}} = 17$	142	62^{+298}_{-44}
(2) $N_{\text{detected}} = 40$	200	135^{+279}_{-86}
(3) $N_{\text{detected}} = 40$, extended prior	200	123^{+412}_{-79}
(4) $N_{\text{detected}} = 40$, $S_{\text{diff}} = 2$ mJy	200	168^{+200}_{-78}
(5) $N_{\text{detected}} = 36$, $p_{\text{fluxless}} = 0\%$	200	124^{+313}_{-86}
(6) $N_{\text{detected}} = 40$, $p_{\text{fluxless}} = 10\%$	200	133^{+277}_{-86}
(7) $N_{\text{detected}} = 34$, $S_{\text{diff}} = 2$ mJy, $p_{\text{fluxless}} = 10\%$	200	222^{+229}_{-128}
(8) $N_{\text{detected}} = 40$, $S_{\text{diff}} = 2$ mJy, $p_{\text{fluxless}} = 10\%$	200	185^{+240}_{-94}
(9) $N_{\text{detected}} = 51$, $S_{\text{diff}} = 2$ mJy, $p_{\text{fluxless}} = 10\%$	200	188^{+201}_{-84}

One may wonder how well the *swyft* posteriors express constraints on the inferred parameters given the information in the data sets. In particular, we are interested in the coverage of the posterior distributions. Using our MSP population simulator, we generate 1000 mock observations of our selected parameter tuple. We infer with our trained network the posterior distributions for each mock observation and count for how many observations the credible intervals from $x\% \in [0, 1]$ encompass the true value. We provide the results of this coverage test in the right panel of Figure 4. This derived coverage demonstrates that our *swyft* setup produces reasonably calibrated posterior distributions for all parameters. Some of the coverage profiles are slightly conservative, especially the inference on $S_{\text{th},\infty}$. However, this parameter does not bear critical information about the MSP population since it is technically known from the sample we use. We emphasize that these coverage tests do not make any statement about how well we have made use of the available information in the data set but rather assess the validity of the

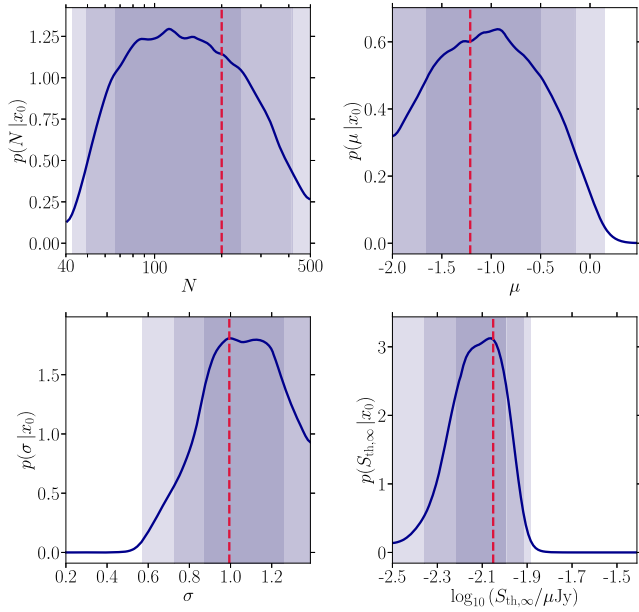


Figure 5. Same as Figure 4 for our baseline target mock observation with parameters $N_{\text{detected}} = 40$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, $S_{\text{th},\infty} = 9 \mu\text{Jy}$, and $p_{\text{fluxless}} = 0\%$.

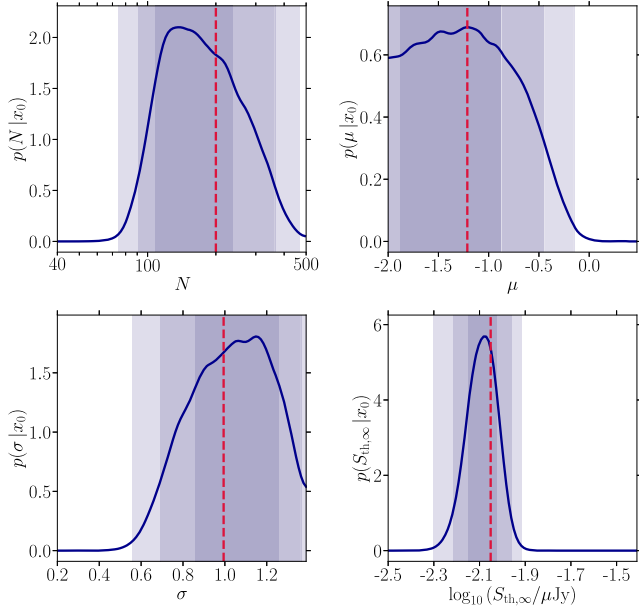


Figure 6. Same as Figure 5 including information about the diffuse radio emission in the simulator for which we assume $S_{\text{diff}} = 2 \text{ mJy}$ and $p_{\text{fluxless}} = 0\%$.

posteriors with respect to the inference problem at hand. The `swyft` coverage results are in stark contrast to those of the Bayesian analysis in Figure 2. While we demonstrated that both methods produce comparable posterior distributions, their statistical calibration differs substantially. We find that the `swyft` pipeline is a statistically more robust approach to infer the properties of MSP populations in GCs.

Yet, the MSP population parameters chosen here do not align well with our current knowledge of the MSP content of Terzan 5 (see Section 2.1.2). To better assess the capabilities of the `swyft` approach we switch to a mock data definition that more closely resembles the real Ter 5 data set. In the following

section, we will describe and examine this baseline setup in greater detail.

5.1.2. Baseline Case

Terzan 5 contains around 40 MSPs with measured fluxes. We adopt this number as N_{detected} in what follows. Because of this large number of known MSPs, we increase the total size of our mock MSP population to $N = 200$ while we keep the luminosity parameters as in the previous section, that is, $\mu = -1.2$ and $\sigma = 1.0$. To achieve the targeted number of detected sources with these parameters, we have to set the detection threshold to $S_{\text{th},\infty} = 9 \mu\text{Jy}$. This value is close to the lowest detected flux of $8 \mu\text{Jy}$ of an MSP in Ter 5 (Padmanabh et al. 2024) at 1284 MHz. Our reported inference results refer to this frequency band, and we notice that our choices for this baseline mock population are reasonable and informed by observations. We start in an ideal setting where all detected MSPs come with a secured flux measurement ($p_{\text{fluxless}} = 0\%$). Finally, we do not include information from the diffuse flux measurement. Our results for this baseline case are illustrated in Figure 5 following the same style as the left panel of Figure 4. The median and the 95% credible intervals of N are reported in the second line of Table 3. We show the coverage results for this case in Figure 10 of Appendix B. All the following findings are accompanied by their coverage reported in Appendix B.

We find that the injected values can be recovered with much better precision than in the Chennamangalam et al. (2013) setup; for all four parameters, we are able to recover well-defined credible intervals and not mere upper limits as for the number of sources N in Figure 4. The detection threshold is reconstructed with the highest precision (and accuracy). This result is very intuitive since the detection threshold is given by the smallest number in the flux column of our synthetic catalogs. We notice, however, that this parameter can already be quite constrained by observations, especially in the case of a uniform survey. More relevant are the MSP population parameters. Here, we find that all injected parameter values are recovered within the 1σ credible interval of the inferred posteriors. Yet, the posteriors still allow for a wide range of MSP population scenarios; in particular, there might be 50–400 MSPs in this mock population at the 2σ level. A similar statement holds for the viable range of μ .

While our posterior distributions do not run against the prior boundaries, it is interesting to check how robust these inference results are against widening the original priors. As N is not strictly constrained in this baseline setup, we launch an alternative `swyft` run with a larger prior from 40 to 3000 sources while keeping the other priors the same. The median and the 95% credible regions of N are reported in the third line of Table 3. The results are shown in the left panel of Figure 16 of Appendix C. Inspection of obtained posterior distributions does not reveal major changes to the results shown here in the main text. They are, in fact, almost identical, so that we proceed by only considering the prior definitions detailed in Table 1.

5.1.3. Impact of Diffuse Radio Emission

To include information about the diffuse radio flux, we assume a value of $S_{\text{diff}} = 2 \text{ mJy}$ that adds on top of the cumulative flux of all detected MSPs, as per Equation (5). The results of the parameter estimation are shown in Figure 6 and the median and the 95% credible intervals of N are reported in

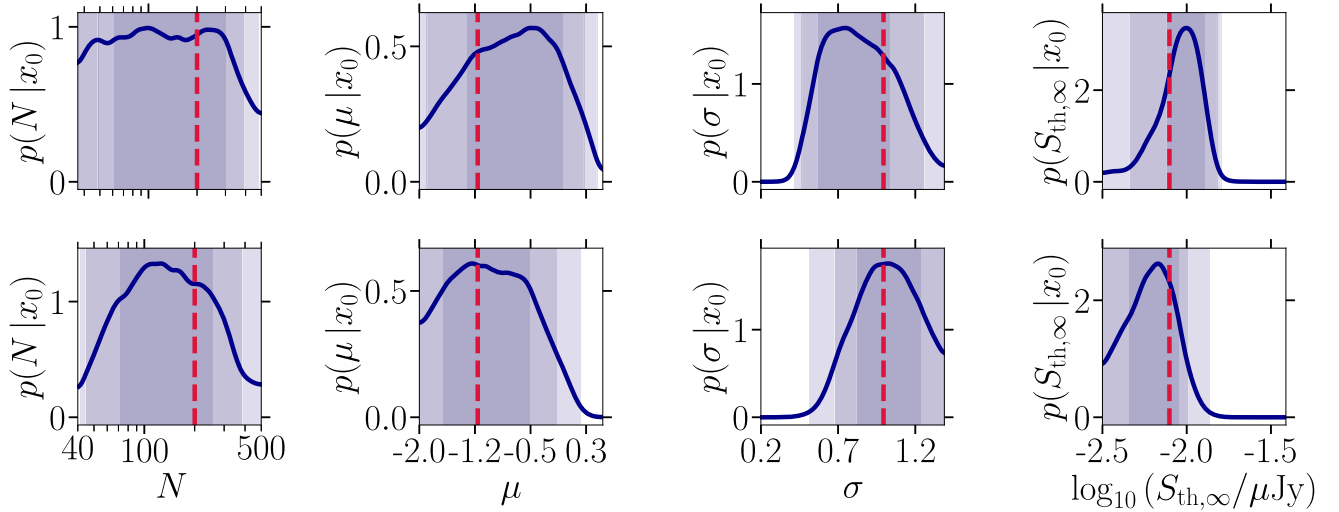


Figure 7. Impact of incomplete MSP flux measurements. Complete set of one-dimensional marginal posterior distributions adopting the color-coding of Figure 4 characterizing the radio MSP population model outlined in Section 4.1 characterized by $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{th},\infty} = 9 \mu\text{Jy}$. We contrast two cases: (upper row) $N_{\text{detected}} = 36$ and $p_{\text{fluxless}} = 0\%$, representing observational results with fewer MSPs than expected regarding the mean expectation of $\langle N_{\text{detected}} \rangle = 40$; (lower row) $N_{\text{detected}} = 40$ and $p_{\text{fluxless}} = 10\%$, illustrating a scenario where $N_{\text{detected}} = \langle N_{\text{detected}} \rangle$ but 10% of the MSPs lack a flux measurement.

the fourth line of Table 3. The corresponding coverage results are shown in Figure 11 of Appendix B.

Compared to our baseline case (Section 5.1.2), adding information about the diffuse emission significantly improves the parameter estimation of the total number of sources N by narrowing the 2σ credible interval to $N \sim [90, 370]$, while we find $N \sim [66, 410]$ without diffuse flux information. We notice a similar striking impact on the precision of the detection threshold’s posterior, which is much more peaked around the injected value. The quality of the inference on the MSP population’s luminosity function remains approximately the same. From a physics perspective, these findings follow from the fact that the assumed diffuse flux is (at least partially) constituted by the subthreshold population, which we can access this way.

5.1.4. Impact of Incomplete Flux Measurements

We investigate the impact of increasing p_{fluxless} as follows: Missing flux information for some detected pulsars implies that the observational source catalog contains fewer sources with complete data than what is expected on average for a given flux threshold $S_{\text{th},\infty}$. We see this situation in the case of Terzan 5. Therefore, we compare in Figure 7 (i) the inferred one-dimensional marginal posterior distributions for a complete catalog of our baseline mock target but assuming fewer detected sources than possible, $N_{\text{detected}} = 36$ (upper row), and (ii) a case with $N_{\text{detected}} = 40$ (lower panel) as expected on average but including four sources, i.e., $p_{\text{fluxless}} = 10\%$, without flux measurements. The corresponding coverage results are shown in Figure 12 of Appendix B, and the median and the 95% interval credible region on N are reported in the fifth and sixth lines of Table 3.

We observe that the quality of the inference overall improves when including detected sources without derived fluxes. While the total number of sources N is rather unconstrained in the upper panel with $p_{\text{fluxless}} = 0\%$, we find a narrower prior in the lower panel. The widths of the 1 and 2σ credible intervals are not very different from Figure 5, where we assume a complete catalog of 40 detected sources. This can be understood from the setup of our training data since we generate a data column that

counts detected sources no matter the availability of a flux measurement. A similar improvement is found for the two parameters of the MSP population’s luminosity function, which are slightly narrower and more centered on the injected value in comparison with the upper row. However, the inferred posterior of the detection threshold deteriorates to some degree as it appears to be wider for the case of $p_{\text{fluxless}} = 10\%$, but the difference is only marginal and, again, comparable to the results shown earlier with full flux information. Therefore, we conclude that adding detected sources irrespective of the availability of an associated flux measurement is beneficial for the ultimate inference.

Finally, we note that, since our method is based on the study of the luminosity function of MSPs in GCs, it would not make sense to apply our method to clusters where a majority of pulsars lack flux measurements. In such cases, one could decide to study the period distribution instead, which is beyond the scope of this paper.

5.1.5. Prospects for Deeper Surveys

Future surveys will attain lower detection thresholds, which will help us to better determine the MSP populations in GCs. To understand how $S_{\text{th},\infty}$ affects the parameter estimation, we assume a survey that measures the diffuse flux in the GC (again $S_{\text{diff}} = 2 \text{ mJy}$ in a scenario where still 10% of the MSPs are reported without measured flux values). We consider three detection thresholds, $S_{\text{th},\infty} = 12 \mu\text{Jy}$ (left column), $S_{\text{th},\infty} = 9 \mu\text{Jy}$ (middle column), and $S_{\text{th},\infty} = 6 \mu\text{Jy}$ (right column), as shown in terms of one-dimensional marginal posteriors in Figure 8 (the corresponding coverage results are shown in Figure 13 of Appendix B), and the median and the 95% credible region of N are reported in lines seven to nine of Table 3. We show results for $S_{\text{th},\infty} = 12 \mu\text{Jy}$ —which is worse than the assumed value of our baseline mock setup—for pedagogical reasons to better illustrate the evolution of the posteriors with decreasing detection threshold. Since a given detection threshold uniquely determines the mean number of detected MSPs, we adjust the assumed value of N_{detected} as described in Section 5.1.1. Consequently, we find $N_{\text{detected}} = 34, 40$, and 51 for the three considered values of $S_{\text{th},\infty}$, respectively.

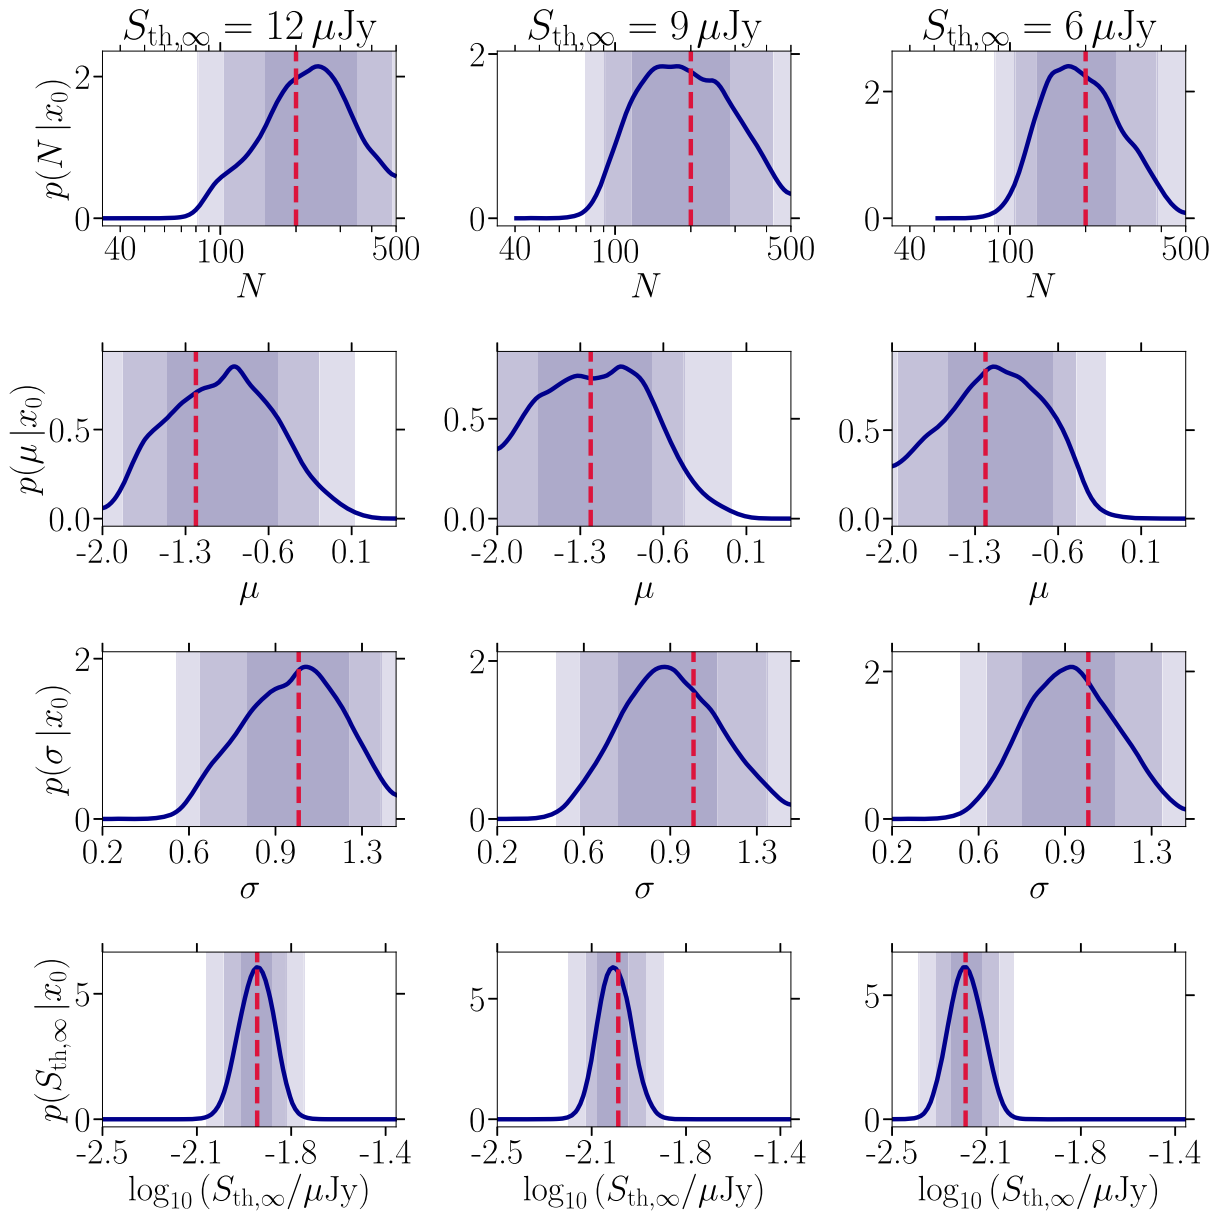


Figure 8. Comparison of the complete set of one-dimensional marginal posterior distributions (blue curves) for the four parameters characterizing the radio MSP population model outlined in Section 4.1. In each panel we show varying minimal detection thresholds for large pulsation periods $S_{\text{th},\infty}$ and accordingly adjusting the expected number of detected sources given the total number of MSPs. The stated values for N_{detected} represent the mean expectation for the respective detection threshold obtained from 1000 realizations. The inference has been performed on target mock observations with shared parameters $p_{\text{fluxless}} = 10\%$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{diff}} = 2 \text{ mJy}$. The color-coding is as in Figure 5. Left column: $N_{\text{detected}} = 34$, $S_{\text{th},\infty} = 12 \text{ } \mu\text{Jy}$. Middle column: $N_{\text{detected}} = 40$, $S_{\text{th},\infty} = 9 \text{ } \mu\text{Jy}$. Right column: $N_{\text{detected}} = 51$, $S_{\text{th},\infty} = 6 \text{ } \mu\text{Jy}$.

The most striking improvement is achieved when lowering the detection threshold from 9 to 6 μJy , while the initial step from 12 to 9 μJy does not have sizeable effects on the final posteriors. In fact, the second step to 6 μJy adds 11 MSPs to the source catalog, while the first step only adds six. We find that the posterior distribution of N becomes much narrower positioning the 2σ credible interval between 100 and 390 MSPs. Improving the detection threshold by a factor of 2 renders this parameter accessible in the posterior estimation with a defined 68% and 90% credibility interval. This is a natural consequence of slightly increasing the number of detected sources because it increases the part of the luminosity function accessible to the inference pipeline. Note that this setup uses 10% of fluxless MSPs as well as the inclusion of a

measurement of the diffuse flux for the mock GC. The inference on the remaining three parameters becomes better but to a lesser extent. For completeness, we also considered a case where all sources are above the threshold, i.e., detected. In this case, the size of the MSP population, the mean and the width of the luminosity function are recovered with small uncertainty. The posteriors and corresponding coverage plots are shown in Figure 14 of Appendix B.

5.2. Real Data: Terzan 5

Having tested the performance of the `swyft` implementation on purely simulated data, we now venture to apply the approach to the available data for Terzan 5 presented in Section 2.1.2. To show the evolution of the posterior

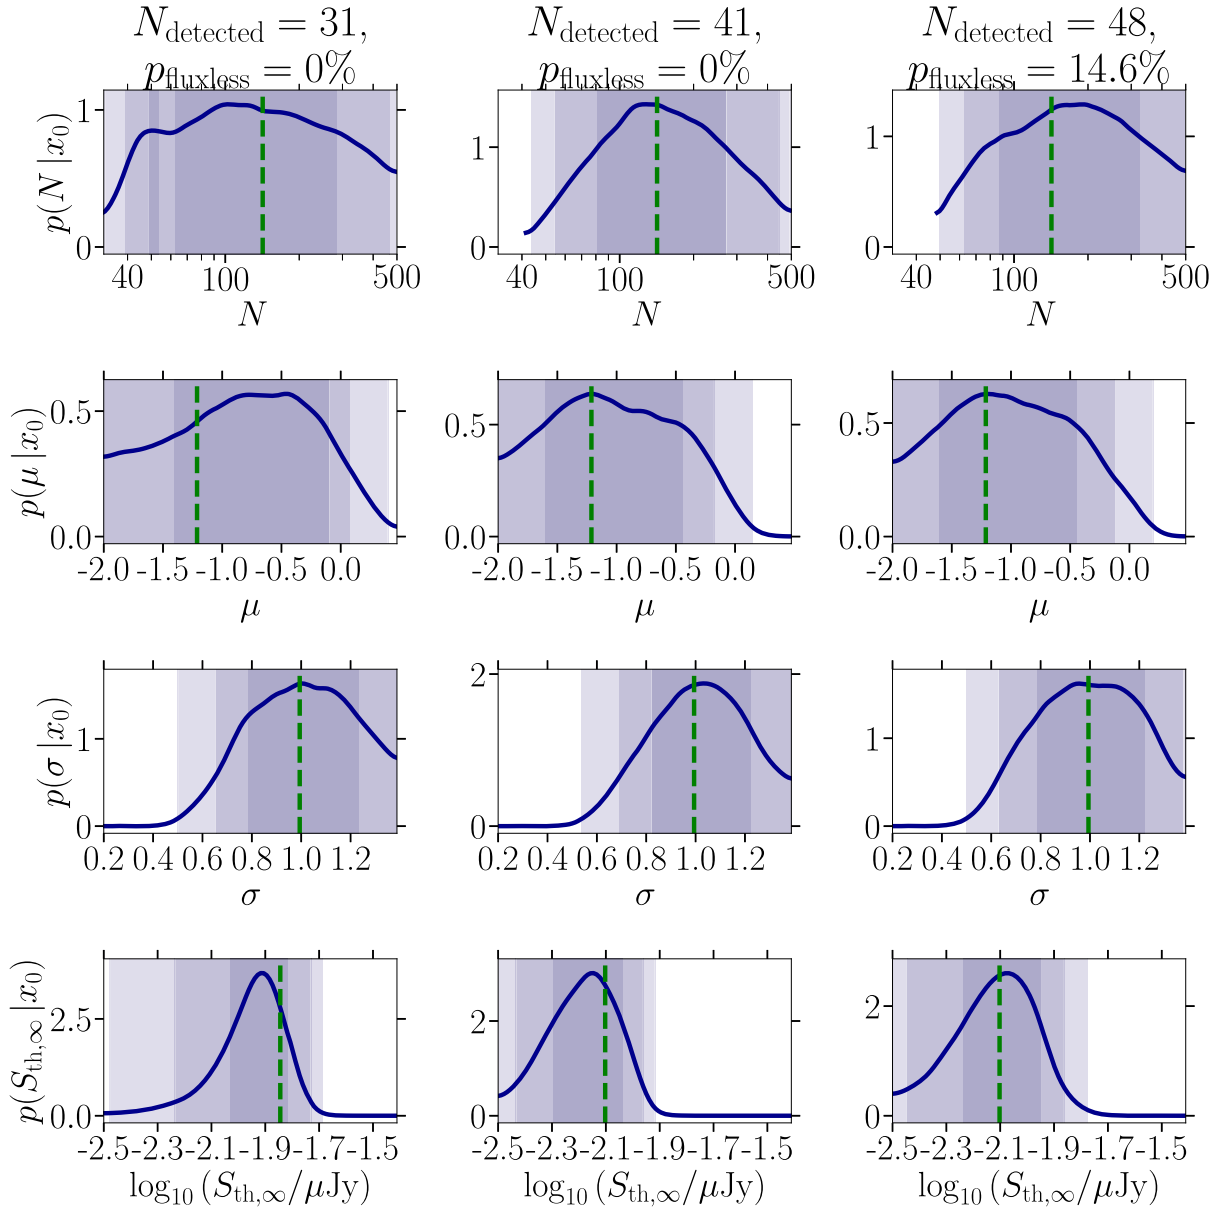


Figure 9. Comparison of the complete set of one-dimensional marginal posterior distributions (blue curves) for the four parameters characterizing the radio MSP population model outlined in Section 4.1 obtained for real MSP catalogs of Terzan 5 outlined in Section 2.1.2. We consider three data sets: (left column) original sample of 31 MSPs with flux measurements (Martsen et al. 2022) rescaled to a frequency of 1284 MHz, (middle column) 41 MSPs with flux measurements encompassing the Martsen et al. (2022) sample and 10 additional sources characterized by MeerKAT at the same frequency, and (right column) adding seven additional sources to the sample without flux measurement ($p_{\text{fluxless}} \approx 14.6\%$) to account for all sources currently known MSPs in Terzan 5. The dashed vertical green lines for N , μ , and σ denote the inference results of Chennamangalam et al. (2013) while the respective line for $S_{\text{th},\infty}$ is the one we used. The inference has been performed on a target mock observation without a diffuse measurement. The color-coding is as in Figure 5.

Table 4

Median and 95% Credible Interval of the Number of MSPs N , the Mean μ , and the Width σ of the Luminosity Function Obtained through the `swyft` Analysis Applied to Real Data for Various Cases

Case	N	μ	σ
$N_{\text{detected}} = 31$	126_{-89}^{+320}	$-0.81_{-1.12}^{+1.01}$	$1.01_{-0.40}^{+0.35}$
$N_{\text{detected}} = 41$	146_{-94}^{+283}	$-1.05_{-0.89}^{+0.99}$	$1.02_{-0.38}^{+0.34}$
$N_{\text{detected}} = 48$	158_{-104}^{+294}	$-1.02_{-0.91}^{+1.02}$	$1.00_{-0.39}^{+0.36}$

distributions with the amount of available data, i.e., an increasing number of detected sources with and without flux measurement, we define three data sets: (i) the 31 MSPs

originally listed by Martsen et al. (2022), (ii) adding the 10 MSPs analyzed by MeerKAT, and (iii) using all 48 known MSPs in Terzan 5, seven of them without flux measurements ($p_{\text{fluxless}} \approx 14.6\%$). To ensure comparability to previous results on mock data, we keep the prior distributions outlined in Table 1. In particular, this implies $N \in [N_{\text{detected}}, 500]$ and $\mu \in [-2.0, 0.5]$. This also allows us to easily compare the `swyft` results with previous work in Chennamangalam et al. (2013) and the application of our Bayesian approach in Section 3.2. The results are displayed in Figure 9 and the median and the 95% credible intervals of N , μ , and σ are reported in Table 4. The corresponding coverage results are shown in Figure 15 of Appendix B.

To facilitate the comparison with the results obtained by the authors of Chennamangalam et al. (2013), we display their inferred median parameters for Terzan 5 as dashed vertical green lines. As concerns the posterior distribution for the total number of MSPs in Terzan 5, we observe that increasing the sample leads to a better definition of the posterior distribution, which evolves from a rather broad posterior in case (i) to a well-defined two-sided posterior in cases (ii) and (iii). This behavior follows our inference findings on the baseline mock MSP population. In data selection (iii), we obtain a 2σ credible interval of $N \in [54, 452]$. We want to emphasize here that with the stated Ter 5 catalogs we are in a different situation than with our mock data sets. It is not clear that the number of detected sources associated with the assumed value of $S_{\text{th},\infty}$ (based on the lowest flux value in the list of detected MSPs) is indeed the average number of detected sources one would expect for the true population parameters of Ter 5. It is much more likely that we are always below the average expectations and much closer to the scenario considered in the upper panel of Figure 7. The quality of our inferred posteriors might suffer from a certain out-of-domain effect regarding simulated training samples and reality. Yet, the coverage of our `swyft` pipeline is well calibrated (see Appendix B) so that the results presented here are the best we are currently able to derive.

With respect to the sizable improvement from (i) to (ii), the evolution of the posteriors for the parameters associated with the luminosity function shows only a mild gain when going from (ii) to (iii). Our posteriors are consistent with the results of Chennamangalam et al. (2013). In this case, the inclusion of fluxless MSPs is not as beneficial as having complete flux information because it renders the inference on the detection threshold less constraining. This effect we could also observe in our study of the impact of p_{fluxless} in Figure 7, where the lower row shows a wider posterior for the detection threshold than the corresponding upper row. We thus argue that obtaining a complete assessment of radio fluxes is essential to pinpoint the properties of the MSP population in Terzan 5 with `swyft`.

6. Discussion

The current limiting factor in our understanding of the population of MSPs in GCs is the observational data. The available MSP catalogs are limited by the performance of radio instruments and by their nonuniformity. Unsurprisingly, our ability to robustly infer the number of MSPs hosted by a cluster will increase as the detection threshold of radio telescopes decreases (Section 5.1.5). Not only an improved detection threshold can improve the inference, but also a more systematic measurement of pulsar fluxes. We showed that the fraction of sources not having a flux measurement is a key element in the inference pipeline both on mock and real data (Sections 5.1.4 and 5.2).

Moreover, as we have seen in Section 5.1.3, incorporating a measurement of the diffuse radio emission can be a crucial element of the parameter estimation. Terzan 5 can be found in the NRAO VLA Sky Survey (NVSS; Condon et al. 1998) catalog as NVSS 174804-244641, a source with a total flux density of $S_{\text{tot}} = 3.4$ mJy at 1.4 GHz. We note that S_{det} , the sum of all MSP fluxes measured in Terzan 5, is already larger than 3.4 mJy. This apparent inconsistency is probably partly due to the fact that the position and the size of NVSS 174804-244641 do not exactly match those of Terzan 5 commonly assumed

today. We thus argue that new measurements of the total radio emission of the GC will highly benefit the analysis of the properties of its MSP population. Nonetheless, one has to keep in mind that, while MSPs seem to be the main contributors to the radio emission of GCs, other populations could contribute, such as low-mass X-ray binaries and accreting stellar-mass black holes (Urquhart et al. 2020). Finally, data of several GCs could be combined, if the luminosity function is assumed to be unique, as done by Bagchi et al. (2011). This will enable the usage of an extended data set, if uniform and accessible. If the detection threshold at large periods $S_{\text{th},\infty}$ does not vary too much from one GC to the next, the luminosity of the dimmest object in each cluster can fluctuate as their distance varies. Therefore, each GC probes a different region of the luminosity function. In the Bayesian analysis framework, each cluster has its own independent likelihood \mathcal{L}_j , and the total likelihood is the product of all these likelihoods (see Appendix E for more details). With `swyft`, only minor modifications of the simulator will instead be necessary.

After the publication of Chennamangalam et al. (2013), stellar data from the Gaia mission made it possible to derive a much more precise distance measurement for Ter 5. While we kept for the sake of comparability the distance of (5.5 ± 0.9) kpc throughout our Bayesian and `swyft` applications, a better distance characterization is given by $d = (6.620 \pm 0.150)$ kpc in Baumgardt & Vasiliev (2021). These two measurements are consistent, but the more recent one puts Ter 5 about 1 kpc farther than we assume in our analyses. Yet, Equation (3) states that a different distance translates to a mere shift of an MSP's luminosity. Therefore, we do not expect strong qualitative differences in the inference results except for the mean μ of the MSP luminosity function. To quantitatively probe the impact of this updated distance to Ter 5, we inferred the posteriors for the case of $N_{\text{detected}} = 32$ within this setting. The results are shown in Figure 17 (plus coverage) of Appendix D. As expected, the inference results are slightly different from the ones visualized in the corresponding plot in Figure 9 except for the results on μ . Here, the posterior's maximum is shifted to larger luminosities, which is expected due to the increased distance of Ter 5 compared to Figure 9. We conclude that our results are robust against a redefinition of the distance to a single GC and all plots shown in the main text retain their validity.

An improvement of our work toward more realism lies in our implementation of the detection threshold. While Chennamangalam et al. (2013) chose to work with an MSP-independent threshold, we have implemented an MSP-dependent threshold to mimic the radiometer equation. However, the variations from one threshold to the next are random and do not depend on specific parameters of the MSP. As noted above, we know that these thresholds strongly depend on the pulsation period of pulsars. Moreover, the luminosity of radio pulsars should also depend on their period and period derivative (see, e.g., Bagchi 2013). Therefore, further studies could benefit from and include a simulation of the pulsation period of MSPs and its link with the radio emission. Theoretical works and simulations could help in that respect to guarantee that the inference is robust and correct.

7. Conclusions

The question of the number of MSPs in GCs has been studied by different groups via their luminosity function, each method producing different (but not incompatible) results. In

this work, we tackled this problem through an analysis framework which quickly developed in the last decade: likelihood-free/simulation-based inference. Our method produces Bayesian posteriors which quantify the uncertainty on various parameters of interest. While the method used by Bagchi et al. (2011) could be interpreted as part of SBI, it does not produce Bayesian posteriors. As for the work of Chennamangalam et al. (2013), based on a Bayesian likelihood, we have shown in Section 3.2 that it likely produces posteriors that are too conservative, indicating a poor calibration of the statistical analysis. We therefore developed a novel analysis pipeline based on *swyft*. We tested our method on simulated data sets and we demonstrated that our analysis consistently produces well-behaved posteriors. In the mock data sets analysis, the best improvements are achieved when the simulated data sets include a measurement of the diffuse radio emission of the GC, or when we assume a detection threshold at large periods of $\sim 6 \mu\text{Jy}$, that is, twice as low as current surveys. Applied to real data of Terzan 5, our analysis robustly hints at a population of more than 80 to less than 350 MSPs (68% credible interval). Despite the well-behaved posteriors, however, large uncertainties remain in the determination of the 95% credible intervals of the number of MSPs. On the other hand, our 95% credible intervals of the width of the luminosity function remain very similar to those of Chennamangalam et al. (2013), while the upper limit on the mean of the luminosity function can be better determined.

Our results call for dedicated campaigns of flux measurements of MSPs in GCs, as well as deep imaging observations associated with a measurement of the GC diffuse radio emission.

Acknowledgments

We acknowledge useful discussion with N. Anau Montel on the *swyft* implementation technical aspects, as well as with K. Dage on the population of stars in Terzan 5. The authors acknowledge support from the Embassy of France in Canada via the 2022 New Research Collaboration Program from the France Canada Research Fund. J.B. and C.E. acknowledge financial support from the Programme National des Hautes Energies of CNRS/INSU with INP and IN2P3, co-funded by CEA and CNES, from the ‘‘Agence Nationale de la Recherche,’’ grant No. ANR-19-CE310005-01 (PI: F. Calore), and the Centre National d’Etudes Spatiales (CNES). The work of J.B. is supported by NASA under grant No. 80GSFC21M0002. The work of C.E. has been supported by the EOSC Future project which is co-funded by the European Union Horizon Programme call INFRAEOSC-03-2020, grant Agreement 101017536. This publication is supported by the European Union’s Horizon Europe research and innovation program under the Marie Skłodowska-Curie Postdoctoral Fellowship Programme, SMASH co-funded under the grant agreement No. 101081355. M.C. acknowledges financial support from the Centre National d’Etudes Spatiales (CNES). D.H. acknowledges funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canada Research Chairs (CRC) program. This work has been completed thanks to the facilities offered by the Univ. Savoie Mont Blanc—CNRS/IN2P3 MUST computing center, especially with GPUs acquired thanks to the Labex Enigmass R&D Booster program.

Appendix A Flux Data Set

In Table 5, we provide the data of the pulsars in Terzan 5 used in this work.

Table 5
Flux Densities of Terzan 5 Pulsars Used in Our *swyft* Analysis

Pulsar	S_{1284} (μJy)	S_{1400} (μJy)	S_{2000} (μJy)	α
A	...	2700	1700	-1.31(15)
C	...	1100	670	-1.4(2)
D	...	71	45	-1.28(14)
E	...	170	110	-1.16(13)
F	...	55	35	-1.218(95)
G	...	24	22	-0.26(11)
H	...	39	24	-1.33(9)
I	...	95	55	-1.53(11)
K	...	66	39	-1.47(7)
L	...	96	43	-2.26(7)
M	...	140	91	-1.14(11)
N	...	150	100	-1.02(11)
O	...	310	160	-1.9(2)
Q	...	56	36	-1.24(14)
R	...	35	17	-2.07(14)
S	...	20	14	-1.09(13)
T	...	26	15	-1.61(9)
U	...	30	12	-2.491(97)
V	...	100	77	-0.86(7)
W	...	54	31	-1.53(14)
X	...	43	24	-1.63(7)
Y	...	37	29	-0.76(12)
Z	...	30	23	-0.7(2)
aa	...	29	20	-1.00(14)
ab	...	45	23	-1.83(12)
ac	...	31	17	-1.67(15)
ae	...	56	50	-0.3(2)
af	...	33	22	-1.19(12)
ag	...	16	9.2	-1.6(2)
ah	...	14	7.1	-1.9(2)
ai	...	33	28	-0.4(2)
ao	12
ap	15
au	12
ax	8
aq	17
ar	44
at	19
as	10
av	15
aw	10

Note. All data from the top section are from Martsen et al. (2022) and all data from the bottom section are from Padmanabh et al. (2024).

Appendix B Collection of Coverage Results for All Scenarios

In this appendix, we collect and provide coverage results (Figures 10–15) for all radio MSP population scenarios we tested and performed inference on. The parameters of the MSP scenario are stated in the respective plot’s caption.

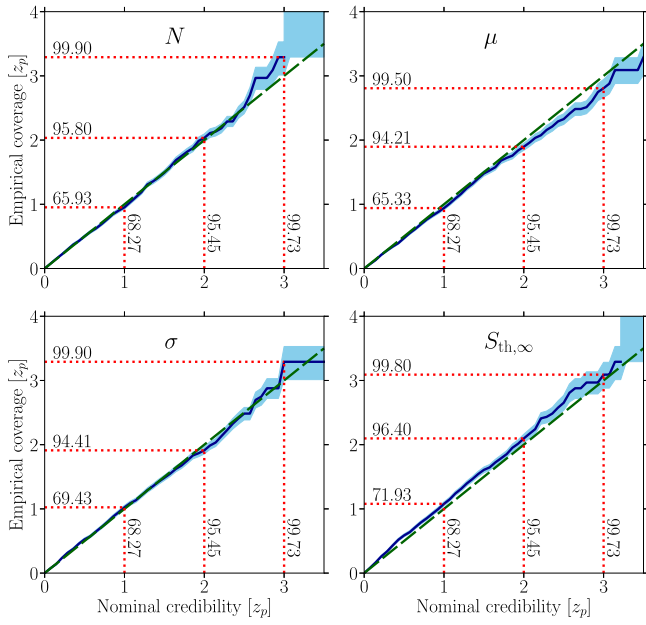


Figure 10. Same as the right panel of Figure 4 in the main text for the baseline mock data setups characterized by $N_{\text{detected}} = 40$, $p_{\text{fluxless}} = 0\%$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{th},\infty} = 9 \mu\text{Jy}$ with a prior range $N \in [N_{\text{detected}}, 500]$ (see Figure 5 for the parameter inference).

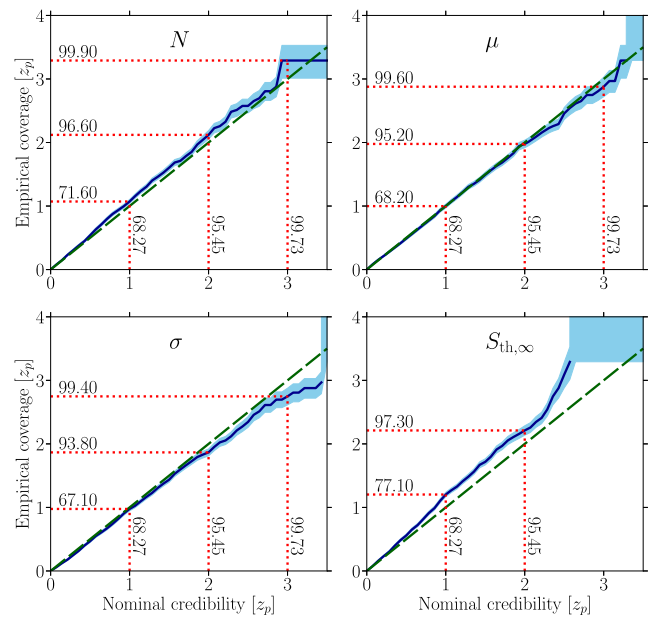


Figure 11. Same as the right panel of Figure 4 in the main text for the baseline mock data setups exploring the impact of adding a diffuse measurement for the GC under scrutiny with results shown in Figure 6: $N_{\text{detected}} = 40$, $p_{\text{fluxless}} = 0\%$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{th},\infty} = 9 \mu\text{Jy}$ and $S_{\text{diff}} = 2 \text{ mJy}$.

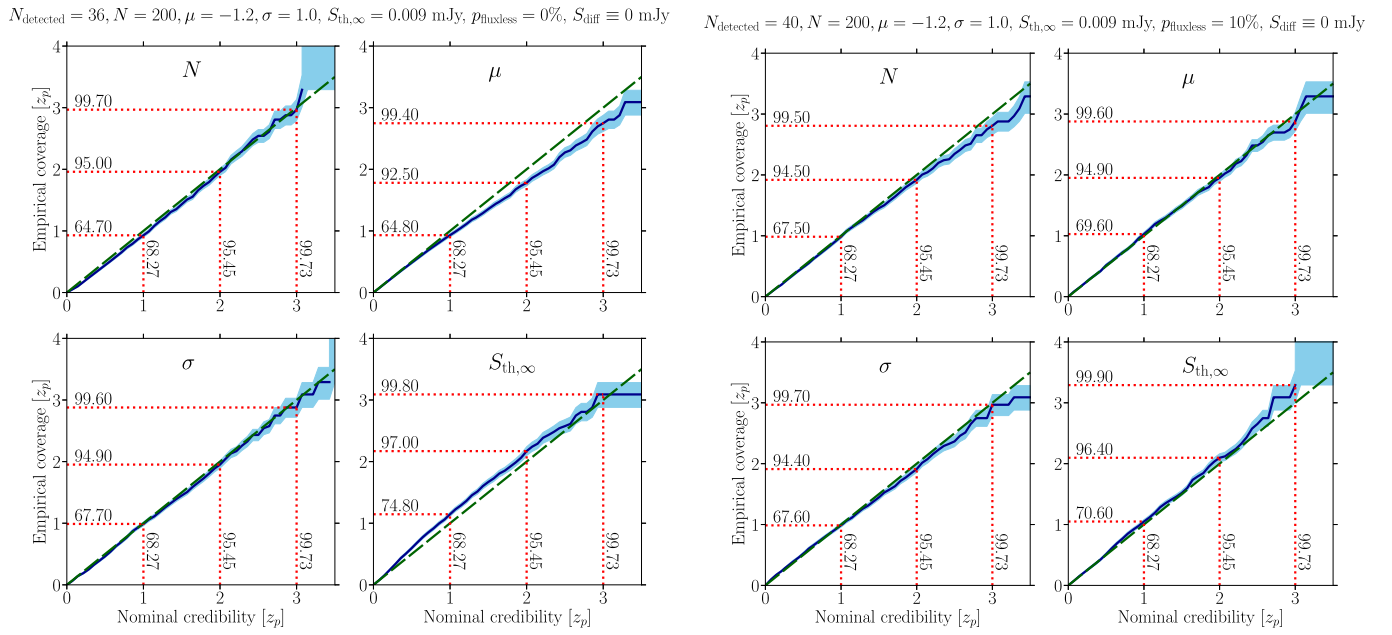


Figure 12. Same as the right panel of Figure 4 in the main text for the baseline mock data setups exploring the impact of having an MSP catalog lacking flux information with results shown in Figure 7: (left) $N_{\text{detected}} = 36, p_{\text{fluxless}} = 0\%, N = 200, \mu = -1.2, \sigma = 1.0,$ and $S_{\text{th},\infty} = 9 \mu\text{Jy}$; (right) $N_{\text{detected}} = 40, p_{\text{fluxless}} = 10\%, N = 200, \mu = -1.2, \sigma = 1.0,$ and $S_{\text{th},\infty} = 9 \mu\text{Jy}$.

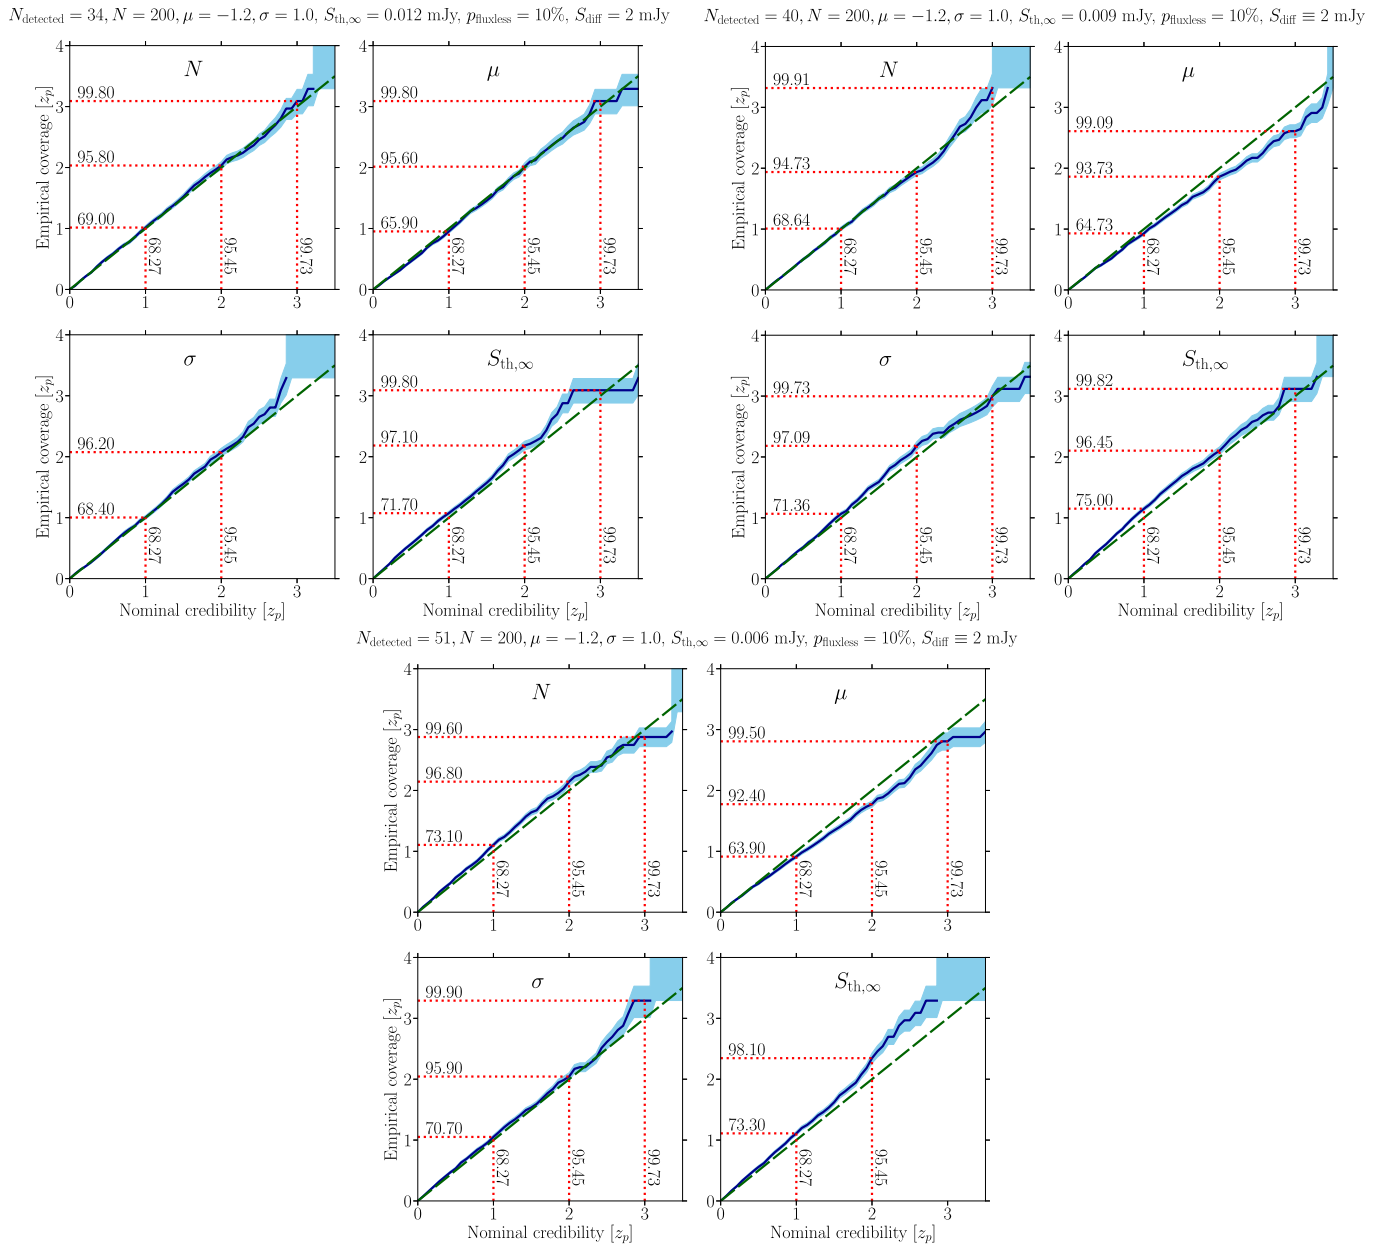


Figure 13. Same as the right panel of Figure 4 in the main text for the baseline mock data setups exploring the impact of lowering the detection threshold with results shown in Figure 8: (upper left) $N_{\text{detected}} = 34$, $p_{\text{fluxless}} = 10\%$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{th},\infty} = 12 \mu\text{Jy}$ and $S_{\text{diff}} = 2 \text{ mJy}$; (upper right) $N_{\text{detected}} = 40$, $p_{\text{fluxless}} = 10\%$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{th},\infty} = 9 \mu\text{Jy}$ and $S_{\text{diff}} = 2 \text{ mJy}$; and (bottom) $N_{\text{detected}} = 51$, $p_{\text{fluxless}} = 10\%$, $N = 200$, $\mu = -1.2$, $\sigma = 1.0$, and $S_{\text{th},\infty} = 6 \mu\text{Jy}$ and $S_{\text{diff}} = 2 \text{ mJy}$.

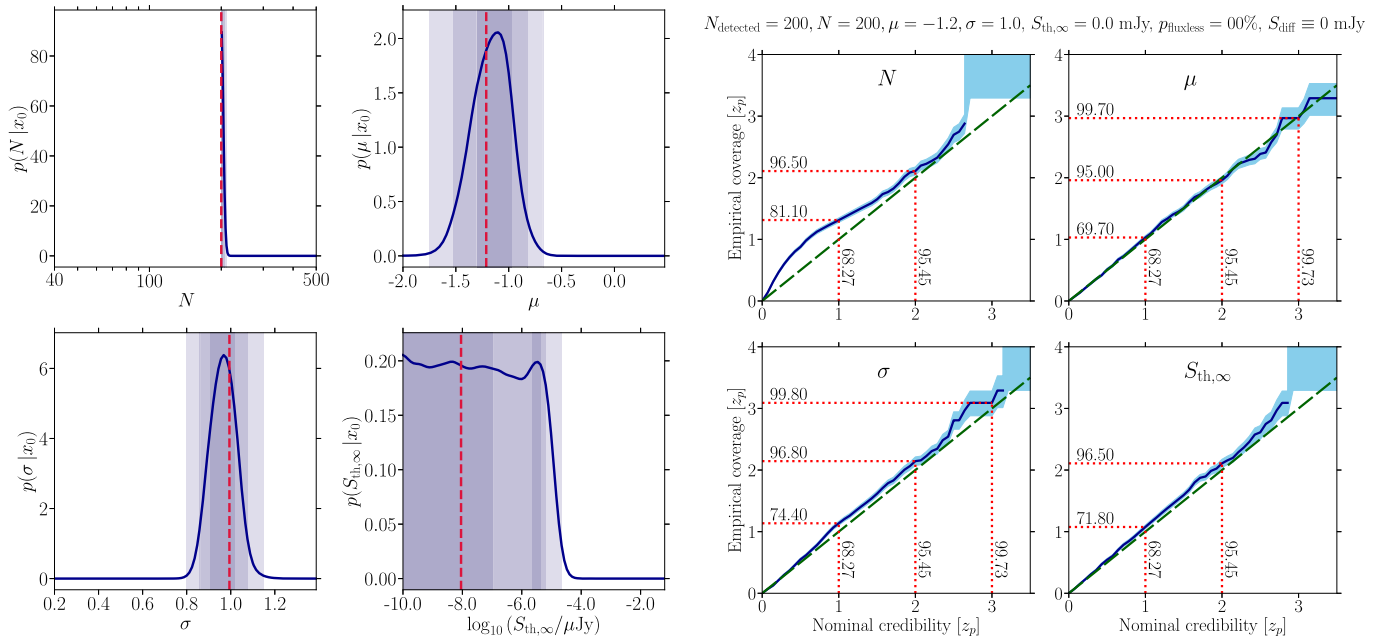


Figure 14. Left: same as Figure 5 considering that all MSPs are above the sensitivity threshold, i.e., $N_{\text{detected}} = N = 200$ and $S_{\text{th},\infty} = 10^{-8} \mu\text{Jy}$. The uncertainty on the mean and width of the radio luminosity function is mostly related to the uncertainty on the distance to Terzan 5 assumed for the mock simulations. Right: coverage plot for scenarios similar to the one considered in the left panel but with $200 \leq N_{\text{detected}} = N \leq 500$. The deviation from the diagonal in the number of sources is due to the sharpness of the posterior. Moreover, being one-sided, it cannot be centered on the true value.

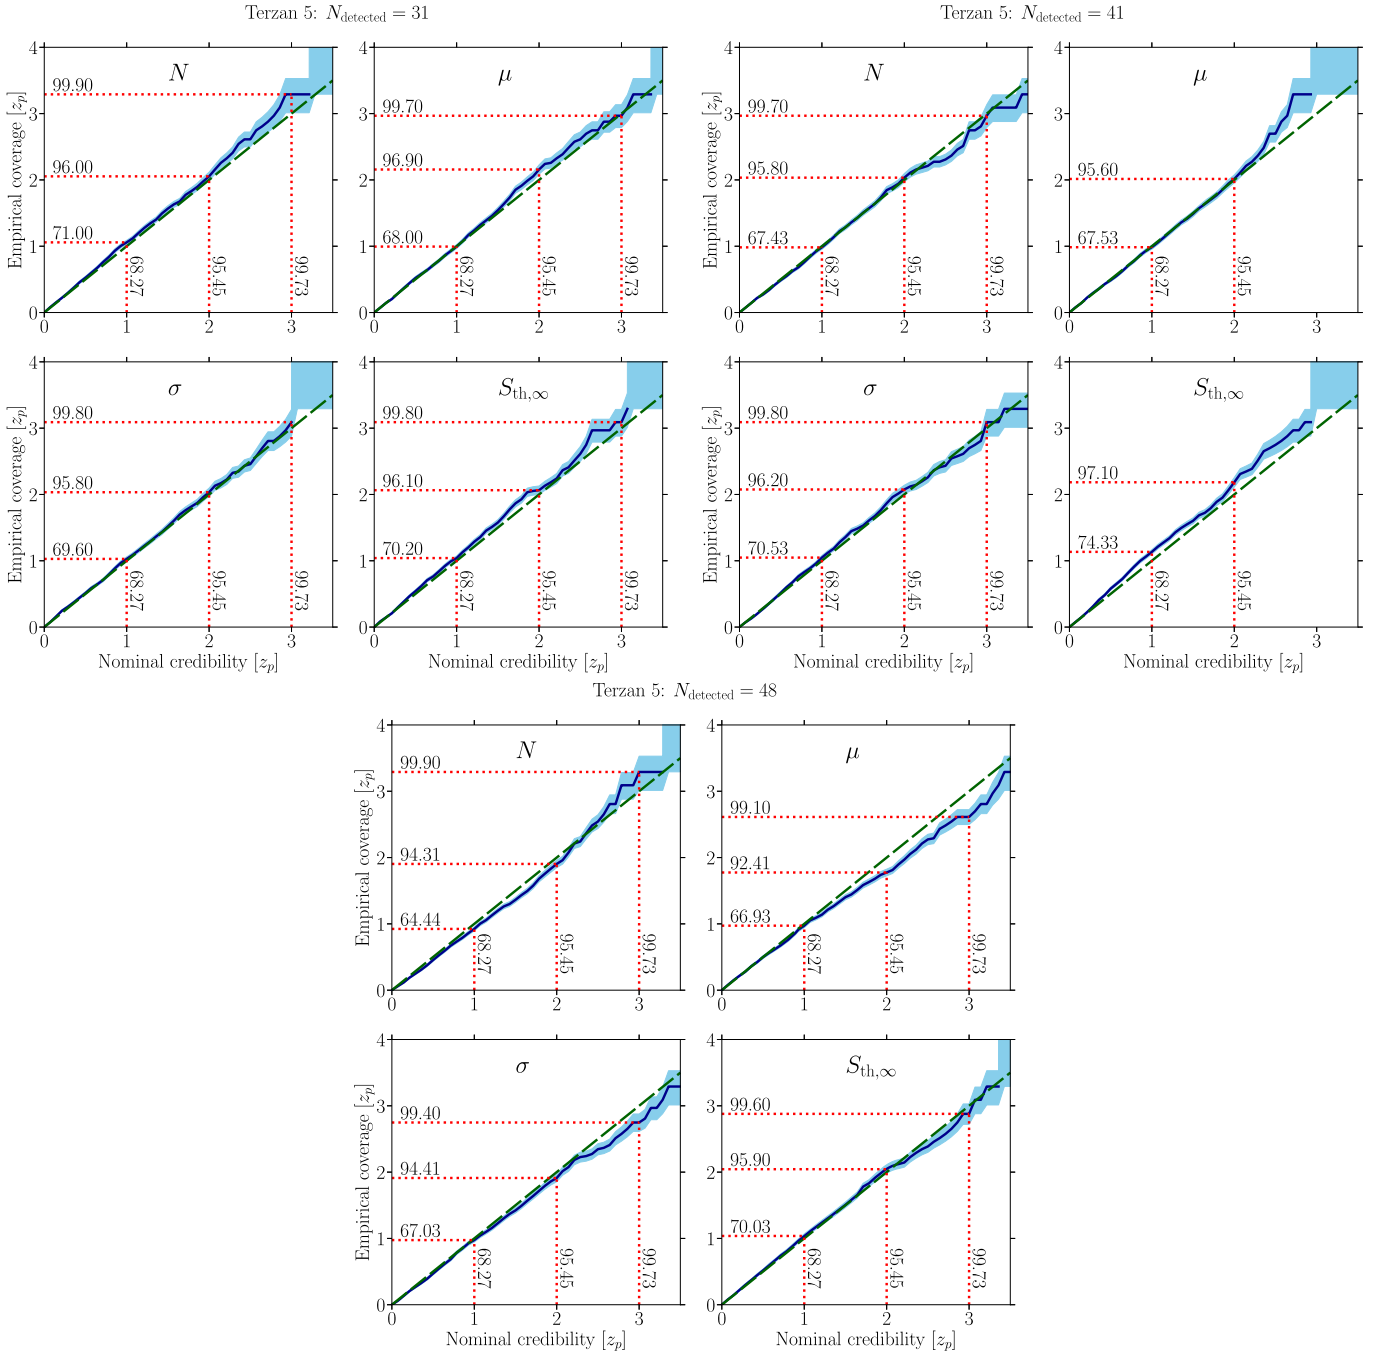


Figure 15. Same as the right panel of Figure 4 in the main text for the radio MSP population of Terzan 5 with (upper left) $N_{\text{detected}} = 32$, $p_{\text{fluxless}} = 0\%$, and $S_{\text{th},\infty} = 16.5 \mu\text{Jy}$; (upper right) $N_{\text{detected}} = 41$, $p_{\text{fluxless}} = 0\%$, and $S_{\text{th},\infty} = 8 \mu\text{Jy}$; and (bottom) $N_{\text{detected}} = 48$, $p_{\text{fluxless}} = 14.6\%$, and $S_{\text{th},\infty} = 8 \mu\text{Jy}$. See Figure 9 for the parameter inference results.

Appendix C

Enlarging the Prior Range on the Total Number of Millisecond Pulsars

In this appendix, we show the results of enlarging the prior range of the `swyft` analysis of our baseline mock data set. We increase the total numbers of sources from $N = 500$ to

$N = 3000$ while keeping the priors of all other parameters identical to the default setup shown in Table 1. The posteriors and the coverage for this case are displayed in Figure 16.

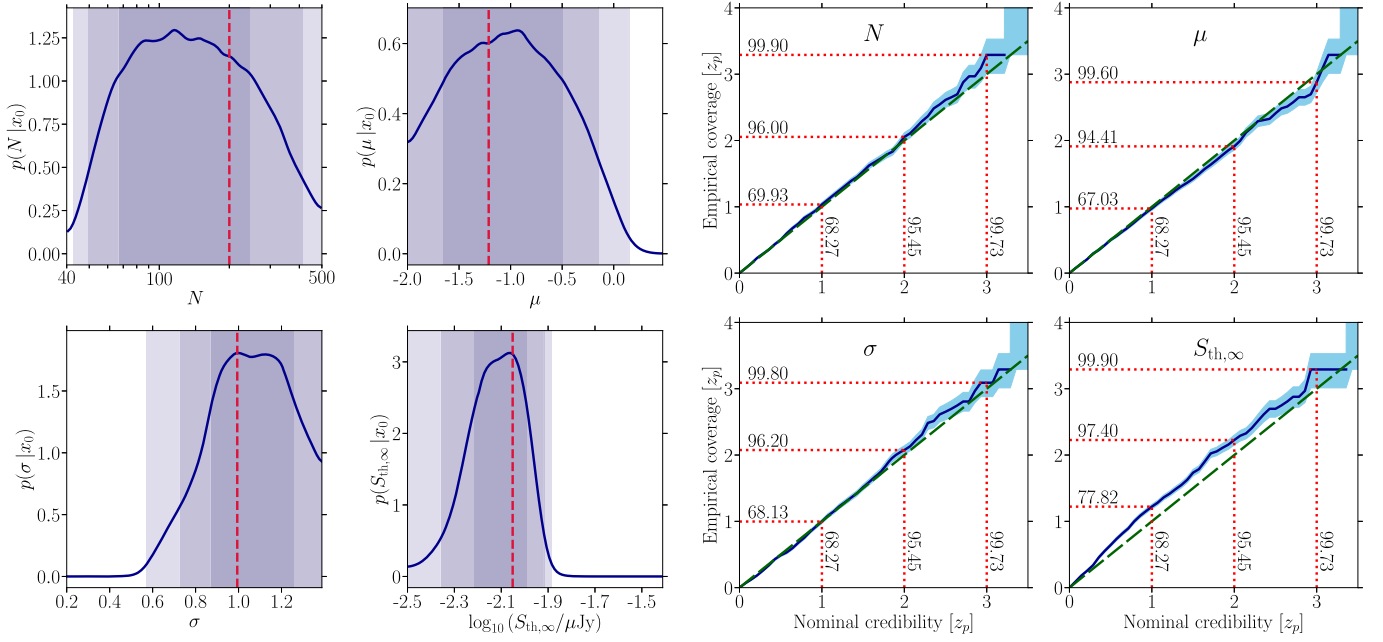


Figure 16. Left: same as Figure 5 using a wider prior on N extending up to 3000 MSPs. Right: coverage plot for the scenario considered in the left panel.

Appendix D Updating the Distance to Terzan 5

In this appendix, we show the results of changing the distance to Ter 5 from 5.5 to 6.62 kpc as inferred from data sets of the Gaia mission in Baumgardt & Vasiliev (2021). We

inspect the case of $n = 32$, that is, the Ter 5 GBT sample, where we show the posteriors and the coverage for this case in Figure 17.

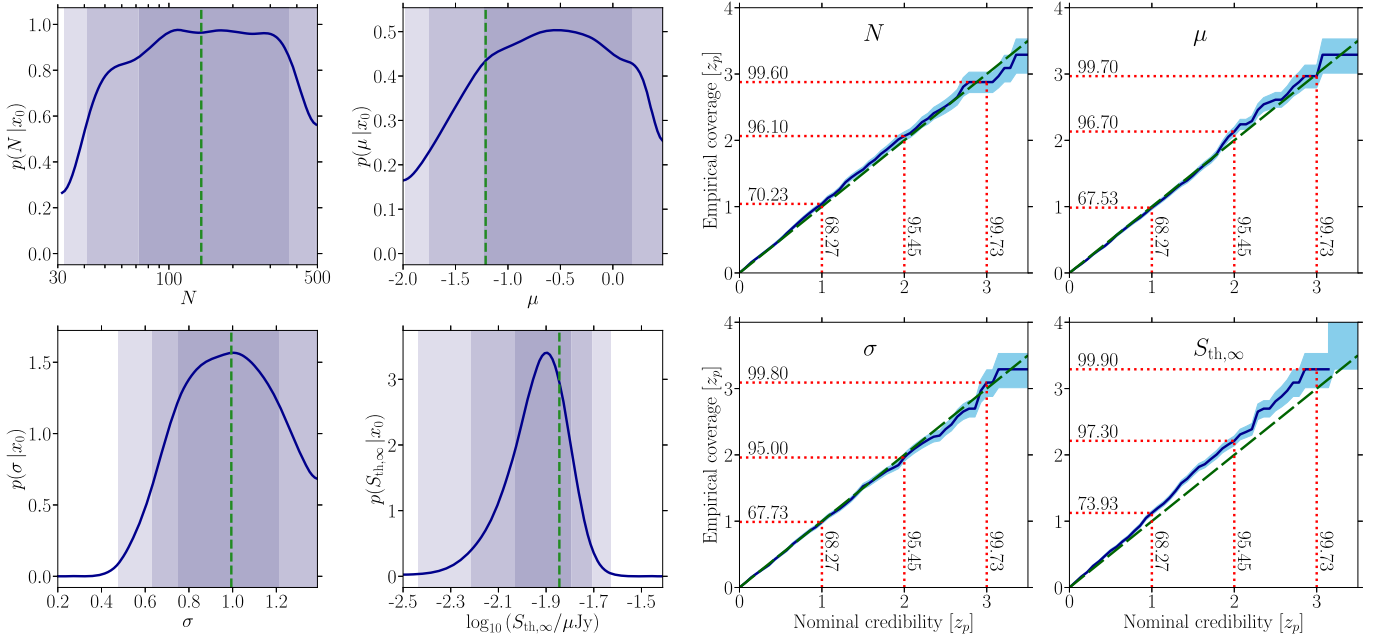


Figure 17. Left: same as Figure 9 using the updated distance to Ter 5 of about 6.62 kpc and the GBT sample of $N_{\text{detected}} = 31$. Right: coverage plot for the scenario considered in the left panel.

Appendix E Multiple Clusters

In Sections 2.2 and 3.1, we introduced several quantities related either to pulsars or to their host cluster, Terzan 5. In order to extend our analysis to several clusters simultaneously, one can use exponents to indicate cluster dependence. As an example, Equation (3) becomes

$$\log_{10}(S_i^j) = \log_{10}(L_i^j) - 2 \log_{10}(d^j). \quad (\text{E1})$$

If the luminosity function is assumed to be the same for all clusters, and therefore that μ and σ are unique, Equation (2) remains valid, and each new cluster only adds three new parameters to the model: N^j , S_{th}^j , and d^j . Otherwise, f , μ , and σ simply become f^j , μ^j , and σ^j , and each new cluster only adds five new parameters to the model. Equation (6) can be used in both cases with $\mathcal{L}(D|\theta, M)$ computed as the product of the $\mathcal{L}^j(D^j|\theta^j, M)$.

ORCID iDs

Joanna Berteaud  <https://orcid.org/0000-0003-4962-145X>
 Christopher Eckner  <https://orcid.org/0000-0002-5135-2909>
 Francesca Calore  <https://orcid.org/0000-0001-7722-6145>
 Maïca Clavel  <https://orcid.org/0000-0003-0724-2742>
 Daryl Haggard  <https://orcid.org/0000-0001-6803-2138>

References

Alvey, J., Bhardwaj, U., Nisanke, S., & Weniger, C. 2023a, arXiv:2308.06318
 Alvey, J., Gerdes, M., & Weniger, C. 2023b, *MNRAS*, **525**, 3662
 Anau Montel, N., Coogan, A., Correa, C., Karchev, K., & Weniger, C. 2023, *MNRAS*, **518**, 2746

Anau Montel, N., & Weniger, C. 2022, in 36th Conf. Neural Information Processing Systems: Workshop on Machine Learning and the Physical Sciences, accepted
 Bagchi, M. 2013, *IJMPD*, **22**, 1330021
 Bagchi, M., Lorimer, D. R., & Chennamangalam, J. 2011, *MNRAS*, **418**, 477
 Baumgardt, H., & Vasiliev, E. 2021, *MNRAS*, **505**, 5957
 Berteaud, J., Eckner, C., Calore, F., Clavel, M., & Haggard, D. 2024, Software Supplementing: Simulation-based Inference of Radio Millisecond Pulsars in Globular Clusters v1, Zenodo, doi:10.5281/zenodo.11403867
 Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *A&A*, **564**, A125
 Chennamangalam, J., Lorimer, D. R., Mandel, I., & Bagchi, M. 2013, *MNRAS*, **431**, 874
 Cole, A., Miller, B. K., Witte, S. J., et al. 2022, *JCAP*, **2022**, 004
 Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, **115**, 1693
 Cranmer, K., Brehmer, J., & Louppe, G. 2020, *PNAS*, **117**, 30055
 Dai, S., Johnston, S., & Hobbs, G. 2017, *MNRAS*, **472**, 1458
 Dewey, R. J., Taylor, J. H., Weisberg, J. M., & Stokes, G. H. 1985, *ApJL*, **294**, L25
 Fruchter, A. S., & Goss, W. M. 2000, *ApJ*, **536**, 865
 He, K., Zhang, X., Ren, S., & Sun, J. 2015, arXiv:1512.03385
 Horangic, D., Orlando, E., Strong, A. W., & Bhattacharyya, S. 2023, in PoS, ECRS (Trieste: SISSA), 117, <https://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=423>
 Martsen, A. R., Ransom, S. M., DeCesar, M. E., et al. 2022, *ApJ*, **941**, 22
 Miller, B. K., Cole, A., Forré, P., Louppe, G., & Weniger, C. 2021, *ANIPS*, **34**, 129
 Miller, B. K., Cole, A., Weniger, C., et al. 2022, *JOSS*, **7**, 4205
 Ortolani, S., Barbuy, B., Bica, E., Zoccali, M., & Renzini, A. 2007, *A&A*, **470**, 1043
 Padmanabh, P. V., Ransom, S. M., Freire, P. C. C., et al. 2024, arXiv:2403.17799
 Pan, Z., Qian, L., Ma, X., et al. 2021, *ApJL*, **915**, L28
 Ridolfi, A., Freire, P. C. C., Gautam, T., et al. 2022, *A&A*, **664**, A27
 Ridolfi, A., Gautam, T., Freire, P. C. C., et al. 2021, *MNRAS*, **504**, 1407
 Sisson, S. A., Fan, Y., & Beaumont, M. A. 2018, arXiv:1802.09720
 Smith, D. A., Abdollahi, S., Ajello, M., et al. 2023, *ApJ*, **958**, 191
 Urquhart, R., Bahramian, A., Strader, J., et al. 2020, *ApJ*, **904**, 147
 Ye, C. S., Kremer, K., Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2019, *ApJ*, **877**, 122