

SPECIAL TRACK: REPRODUCIBLE RESEARCH

Reproducing the Results for Neutron Star Interior Composition Explorer Observation of PSR J0030 + 0451

Chaitanya Afle  and Patrick R. Miles , Syracuse University, Syracuse, NY, 13244, USA

Silvina Caíno-Lores , University of Tennessee, Knoxville, TN, 37996, USA

Collin D. Capano , Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476, Potsdam, Germany

Ingo Tews , Los Alamos National Laboratory, Los Alamos, NM, 87545, USA

Karan Vahi  and Ewa Deelman , University of Southern California, Los Angeles, CA, 90007, USA

Michela Taufer , University of Tennessee, Knoxville, TN, 37996, USA

Duncan A. Brown , Syracuse University, Syracuse, NY, 13244, USA

NASA's Neutron Star Interior Composition Explorer (NICER) observed X-ray emission from the pulsar PSR J0030 + 0451 in 2018. Riley et al. reported Bayesian parameter measurements of the mass and the star's radius using pulse-profile modeling of the X-ray data. This article reproduces their result using the open source software X-Ray Pulse Simulation and Inference and publicly available data within expected statistical errors. We note the challenges we faced in reproducing the results and demonstrate that the analysis can be reproduced and reused in future works by changing the prior distribution for the radius and the sampler configuration. We find no significant change in the measurement of the mass and radius, demonstrating that the original result is robust to these changes. Finally, we provide a containerized working environment that facilitates third-party reproduction of the measurements of the mass and radius of PSR J0030 + 0451 using the NICER observations.

The reproducibility of research—the ability to arrive at a consistent result given the same raw data and original analysis method—is a critical element of scientific discovery. Reproducibility provides the necessary level of trust in the published results and enables researchers to build upon that work. Since more and more scientific studies are using

computation as a tool, reproducibility challenges arise from the computational point of view—especially the availability of data, software, the needed execution environment, and tools as well as documentation used in the original analysis.¹

Neutron Star Interior Composition Explorer (NICER) is a payload onboard the International Space Station, and the X-ray Timing Instrument is dedicated to observing X-rays from galactic pulsars.² Based on NASA's open science and open data policy, the data observed by NICER are released to the public to advance scientific research. One of NICER's aims is to measure the masses and radii of neutron stars. These measurements constrain the neutron star equation of

© 2024 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License. For more information, see <https://creativecommons.org/licenses/by-nc-nd/4.0/>
Digital Object Identifier 10.1109/MCSE.2024.3381080
Date of publication 27 March 2024; date of current version 14 June 2024.

state—the relation between the pressure and density of the neutron star. Measuring this equation of state requires a computationally intensive analysis of the NICER data. To fully understand and leverage the results of the equation of state analyses, the astrophysics community needs to be able to reproduce and modify the original results so that they can 1) check the robustness of the original result, 2) build new analyses using the original result, or 3) extend the original analysis to address new and different questions.

To measure the equation of state using PSR J0030 + 0451, Riley³ developed the analysis software X-Ray Pulse Simulation and Inference (X-PSI).^a X-PSI includes a Bayesian analysis framework to measure the pulsar's mass and equatorial radius (hence, infer the equation of state) using the observed NICER data. We explore whether the analysis of the pulsar PSR J0030 + 0451 by Riley et al.⁴ can be reproduced and modified to test the robustness of the result.

Miller et al.⁵ used the same NICER observations of PSR J0030 + 0451 to produce an independent analysis using different software, models, and methods to measure the mass and radius of the pulsar. They arrived at measurements of mass and radius of PSR J0030 + 0451 that were slightly different from but consistent with the results of Riley et al.⁴ This is an example of the *replicability* of research: using the same data but different methods to arrive at a consistent result. The conclusion can be drawn that the results are replicable; an analysis of the data leads to a consistent mass and radius result for the pulsar. However, it does not verify that an external entity could use the existing software stack created by Riley et al. to achieve the same result, nor does it demonstrate that another group could modify or extend this analysis.

Unlike our previous work on reproducing the detection of GW150914 by Laser Interferometer Gravitational-Wave Observatory,⁶ none of the authors of this reproducibility effort were involved in the original analysis. This work is entirely based on the papers and articles, data, software, and documentation provided to the public by the authors of the original study by Riley et al.⁴ First, we reproduce the results in Figure 19 of Riley et al., which shows the measurement of the mass and the radius of the target pulsar obtained from the analysis. We note the lessons learned and challenges we faced during the reproducibility process, as was done in our previous works where we reproduced the images of the M87 black hole published by the Event Horizon Telescope (EHT) collaboration^{7,8} and reproduced the detection of GW150914.⁶ We discuss

the challenges encountered while acquiring the input data, installing and using the software (including setting up the required dependencies and environment), writing configuration files and job submission scripts, and postprocessing the job output. Ultimately, we were successfully able to reproduce the measurements done in the original analysis.

Going beyond our previous work, after reproducing the original analysis, we demonstrate that Riley et al.⁴ provide sufficient information to allow a third party to modify the analysis in the new work. We use this functionality to test the robustness of the methods to the prior probability distributions chosen for the Bayesian analysis. Specifically, we expand the previous space on the pulsar radius from 16 to 25 km and change the sampler configuration for the Bayesian analysis. We find that changing the upper limit of the prior does not change the posterior distribution statistically significantly, demonstrating the result's robustness to the prior radius choice. We increase the number of sampler points used to sample the posterior probability space from 1000 to 4000 and find that the posterior probability distribution does not change, demonstrating the robustness of the analysis.

As part of our work, we repackage X-PSI and its software dependencies into a Docker container. This aids in the portability of the data and the software and streamlines the reproduction of the original analysis. The container is fully documented and contains the scripts for the entire workflow used in our reanalysis. Scripts to reproduce the container are available from a GitHub repository indexed by Zenodo.^b

This article is organized as follows. First, we describe the original analysis and provide background information on measuring the mass and radius of neutron stars from X-ray data. Then, we describe our effort to reproduce the study and note the computational challenges. Finally, we summarize the lessons we learned when reproducing the Riley et al. analysis and provide guidelines for improvement of the reproducibility of such computationally intensive analyses.

ANALYSIS OF PSR J0030 + 0451

Figure 1 shows a schematic overview of the analysis by Riley et al.,⁴ including the observation of the X-rays by NICER, modeling of the X-ray emission from the pulsar surface, and estimation of the mass and radius of the pulsar using the observational data and the models. The parts of analysis performed by Riley et al. using X-PSI are shown as green boxes.

^a<https://github.com/xpsi-group/xpsi.git>

^b<https://doi.org/10.5281/zenodo.10564150>




FIGURE 1. Schematic showing how the mass and radius of PSR J0030 + 0451 can be measured using X-ray data observed by NICER. The parts of the analysis done by Riley et al.⁴ using X-PSI are shown as green boxes. X-PSI uses the data observed by NICER (box 1 and the observed pulse profile) and the hot spot models that simulate the X-ray emission from the pulsar (box 2) to perform Bayesian parameter estimation (box 3) and measure the posterior probabilities for the mass and radius of the target pulsar (box 4). The two examples of hot spot models shown in the figure are ST+PST and ST+EST. NICER: Neutron Star Interior Composition Explorer; ST+EST: single temperature + eccentric single temperature; ST+PST: single temperature + protruding single temperature; X-PSI: X-Ray Pulse Simulation and Inference; GR: general relativity.

The mass and radius of the neutron star are imprinted on the X-rays emitted by hot spots on the neutron star's surface through the relativistic effect of their propagation through the space-time curvature induced by the star. The X-ray pulse profile detected by a distant observer encodes the neutron star's compactness, the ratio of the star's mass to its radius. NICER measures X-ray counts as a function of time for a target pulsar, as illustrated in box 1 of Figure 1. Since the photon count profile of the signal is identical for each rotation of the pulsar, the signal can be phase-folded into a single pulse profile, which gives the

photon count as a function of the phase of the rotation of the pulsar, as shown in Figure 1 at the top right. Creating a phase-folded dataset is a preprocessing step performed by the NICER instrument team and creates a derived dataset used in subsequent analyses. This dataset, released using Zenodo, is the starting point for the Riley et al.⁴ analyses.

To measure the mass and radius of the star, a model H is created that describes the X-ray emission from the hot spots and uses relativistic ray tracing of the emitted radiation to predict the pulse profile observed by a distant observer. The parameters of this

model are represented by $\vec{\vartheta}$ and include the mass and radius of the neutron star, the parameters describing the geometry of the hot spots, the distance to the pulsar, and the inclination angle of the axis of rotation to the line of view. In box 2 of Figure 1, we show the geometry of the hot spots assumed for two models that we use for reproducing the results from the original analysis.

For a given model H , Bayes' theorem is used to infer the posterior probability distribution of the model parameters given a realization of the observed data (box 3 in Figure 1) according to

$$p(\vec{\vartheta} | \vec{d}, H) = \frac{p(\vec{d} | \vec{\vartheta}, H) p(\vec{\vartheta} | H)}{p(\vec{d} | H)}. \quad (1)$$

Here, $p(\vec{\vartheta} | H)$ is the probability density of the parameters based on theory, assumptions, or previous observations (the prior); $p(\vec{d} | \vec{\vartheta}, H)$ is the joint probability distribution as a function of parameters given fixed data (the likelihood); and $p(\vec{d} | H)$ is the marginalized likelihood, also called evidence.

Riley et al.⁴ aim to produce posterior probability measurements of the mass, radius, and other model parameters for a given pulse-profile model and a set of NICER observations of the pulsar J0030 + 0451. Box 4 in Figure 1 shows an example schematic of a 2-D marginalized posterior of mass and radius. Each parameter's marginalized posterior probability distribution is obtained using the MULTINEST⁹ implementation of the nested sampling algorithm.

Riley et al.⁴ explored several hot spot geometry models to determine which model was most favored. However, we only consider the two most likely used models from their analysis and focus on the posterior probabilities for the parameters of these models. To keep the scope of our work reasonable, we neglect models with lower evidence values or higher complexity.

The hot spot geometry used in the most favored model of Riley et al.⁴ involves two hot regions on the pulsar's surface. The first hot spot is a hot circular disk, whereas the higher temperature in the second hot region lies in the arc-shaped region. This model is named "ST+PST" (for single temperature + protruding single temperature), shown in Figure 1. Using this model, Riley et al. found that the mass of the pulsar PSR J0030 + 0451 is $M = 1.34^{+0.15}_{-0.16} M_{\odot}$, (where $1 M_{\odot}$ is the mass of the sun) and that the equatorial radius is $R_{\text{eq}} = 12.71^{+1.14}_{-1.19}$ km. The bounds mentioned here are the 16% and 84% quantiles from the posterior distribution obtained. The compactness M/R_{eq} is measured to be $0.16^{+0.01}_{-0.01}$. In comparison, Miller et al.⁵ measured the mass and radius to be $M = 1.44^{+0.15}_{-0.14} M_{\odot}$ and $R_{\text{eq}} = 13.02^{+1.24}_{-1.06}$ km. The second Riley et al. model we

investigate is "ST+EST" (single temperature + eccentric single temperature), which differs from ST+PST in that the second hot region is an eccentric annular ring. The model ST+EST gives larger radius and mass ($R_{\text{eq}} = 13.89^{+1.14}_{-1.30}$ km and $M = 1.46^{+0.17}_{-0.18} M_{\odot}$, respectively).

COMPUTATIONAL CONSIDERATIONS

The X-PSI code used by Riley et al.⁴ is an open source code written primarily in Python 2.7 with additional Cython support. As noted in our previous work, Python code presents challenges in reproducibility due to its need for libraries that may not be installed (or may be installed at a different version) on the platform where the code is executed to reproduce an analysis. To address this, Riley et al. provided a Python 2.7 Conda environment to install the code and its dependencies. Although this does not isolate the code in the same way as containerization, it facilitates the installation of X-PSI and necessary libraries at the correct version.

The documentation provided by Riley et al.⁴ indicated that they used version 0.1 (v0.1) of X-PSI to infer the properties of PSR J0030 + 0451; this tagged code was made available on GitHub and was straightforward to obtain and install. Riley et al. also released a Zenodo repository^c that contains the phase-folded X-ray data from NICER used as input to the analysis, the configuration files for X-PSI v0.1, the submit scripts for the job, the output files of the job, and the files containing the posterior samples for each analysis. This thorough release of data and configuration files makes it possible to reproduce the original analysis, given sufficient computational resources. The repository comes with a README.txt that briefly describes each file and its use. For our analysis, we use the configuration files from the repository, changing the paths to the input data and output files wherever necessary. Following the methods of Riley et al., we start our analysis from the phase-folded data and do not attempt to reproduce this from the raw NICER data.

The X-PSI analysis is computationally expensive. For example, obtaining marginalized posteriors for the ST+PST Bayesian analysis took 42,453 wall clock hours (see Table 2 from Riley et al.⁴). The analysis involves $\mathcal{O}(10^8)$ likelihood evaluations, each taking $\mathcal{O}(1)$ s of evaluation time. The likelihood calculation involves simulating hot spots on the star's surface, ray tracing the radiation to include relativistic effects, and creating an instance of the X-ray pulse that a distant observer would detect, making likelihood evaluation the most expensive step. While running X-PSI on a single

^c<https://doi.org/10.5281/zenodo.5506838>

compute node is straightforward, the likelihood evaluations must be executed on multiple compute nodes in parallel to complete the analysis within a reasonable time. To execute the analysis in parallel across multiple compute nodes, X-PSI uses the Message Passing Interface (MPI) library.

To distribute the software and input data on each compute node, we used the Conseil Européen pour la Recherche Nucléaire (CERN) Virtual Machine File System. This shared file system makes scientific software easily available and accessible on high-performance computing (HPC) clusters. The software stack for interprocess communication uses the `mpi4py` library¹⁰ to create Python bindings to MPI libraries written in C++ and installed as dynamic shared libraries. The object code in these libraries executes the interprocess communication using system calls. Our main reproducibility challenge was to produce a containerized version of X-PSI that could execute the analysis using MPI across multiple (possibly heterogeneous) compute nodes.

Following our previous experience, we created a Docker image containing the complete X-PSI software stack for execution as a stand-alone image. Using a base Debian Miniconda image,^d we installed X-PSI and the required dependencies in it using the files provided by Riley et al.⁴ While the Docker container streamlines the installation of X-PSI and its dependencies, it presents a problem of running the analysis in parallel on multiple compute nodes, as code running each Docker image on a compute cluster needs to communicate with the other images. While this is possible, it is challenging without administrative control of the host machines.

Singularity provides a controlled, containerized environment with the advantage that codes running in the image can access the network capabilities of the host. Unlike Docker's full virtual machine containerization, Singularity creates images that overlay on the host machine. Therefore, if the host's operating system is configured to allow interprocess communication for MPI (as is common in cluster environments), it can be used by code running in the Singularity image. For this to work, the Singularity image must contain the exact version of the `mpi4py` and the MPI-shared libraries as the host machine. To address this challenge, we use scripts by the Open Science Grid team that converted Docker containers into Singularity images.

We use HTCondor as the job scheduler and the Syracuse University Gravitational-Wave Group (SUGWG) cluster for our analysis. This is a heterogeneous combination of Intel Xeon Gold 6248 R at 3.00 GHz, E5-2660v2 at 2.20 GHz, E5-2698v3 at 2.30 GHz, X5650 at 2.67 GHz,

X5550 at 2.67 GHz, and E5-2620 0 at 2.00 GHz and AMD EPYC 7702 P and EPYC 7543 processors). The cluster uses a CentOS operating system configured to allow codes to use the OpenMPI¹¹ implementation of MPI.

Since the Docker container (and, hence, the derived Singularity image) used to host X-PSI has a Debian operating system, we made a copy of the OpenMPI shared libraries that are installed on the SUGWG compute nodes, deployed it in the container, and configured the runtime linker so that the Python interpreter inside the container could access these libraries. This ensures that the OpenMPI within the container has paths and configurations identical to the host compute nodes. The analysis can then be launched using the HTCondor job scheduler. HTCondor uses `mpirun` to execute X-PSI from inside the Singularity container across multiple compute nodes. Figure 2 shows the software setup schematic.

REPRODUCING THE J0030 RESULT

We reproduce the original analyses using the ST+PST and ST+EST models using the described software setup. Additionally, we submit jobs using the ST+PST and ST+EST models with broader radius priors than the original analyses. All analyses were run with a sampling efficiency of 0.3 and an evidence tolerance of 0.1, as used by Riley et al.⁴ Table 1 shows the job statistics and summary of results obtained for the analysis by Riley et al. and for our work in this article. We also show the information provided in the Zenodo release by Riley et al. The results from the data in the Zenodo repository and those in the publication are different because Riley et al. postprocessed the posterior samples (present in the Zenodo repository) to reevaluate the evidence. We do not repeat this postprocessing step since we do not have access to a working postprocessing script used by the original study's authors.

Figure 3 shows the posteriors for the model ST+PST from the Zenodo repository (results used by Riley et al.⁴) in blue, from the reproducibility analysis in orange, and for the analysis with the broader radius priors in green. Figure 4 shows when the model ST+EST was used. The figures show 1- and 2-D marginal posteriors for the mass M (in solar masses), the equatorial radius R_{eq} (in kilometers), and the compactness M/R_{eq} (in solar mass/kilometer) of the pulsar. The original and reanalysis analyses involved 19 parameters (for ST+PST and ST+EST models). For brevity, we show the posteriors for only three parameters: mass, radius, and compactness.

We reproduced the result from the original analyses. We get the exact measurement for the three quantities of interest with the same 68 percentile

^d<https://hub.docker.com/r/continuumio/miniconda>




FIGURE 2. Schematic showing the software components used in our workflow execution. We create a Docker image on the login node on the sugwg-condor cluster, which has the replica of the OpenMPI installation present on any of the compute nodes. This Docker image is then pushed to Docker Hub cloud storage. CVMFS converts the Docker container into a Singularity container and makes it available for use on the cluster. We use the HTCondor job scheduler to deploy the Singularity containers on the compute nodes, thus fulfilling the parallelization requirements of the analysis. CVMFS: CERN Virtual Machine File System; MPI: Message Passing Interface.

confidence interval. The minor differences in the posteriors are statistical, and one expects this order of fluctuation after each repetition of the analysis. Since nested sampling is a Monte Carlo technique, one cannot obtain identical reproduction of the positions of the sample points. The samples exploring the posterior space accumulate around the region with high probability. There would be some fluctuation at the periphery of this distribution, which is reflected in the deviations in the 99% contour lines, where the sampler population density is sparse. The inner 68% and 95% contour lines in the posterior from the reanalysis show much less deviation from the posterior of the original analysis.

Although the value for evidence we get for the reanalysis is close to the value reported in the original article, the difference is larger than would be expected from the evidence tolerance provided to MultiNest. Although statistical fluctuations are expected in the evidence, the values obtained illustrate the challenges in comparing evidence between models even when the physical posterior parameters are in good agreement.

Postprocessing and plotting the output of the Bayesian analysis proved to be an obstacle to the reproducibility of the original results. Although

the Zenodo repository had all of the configuration files and submitted scripts to start the analysis, the postprocessing and plotting scripts were absent. These scripts are necessary to produce a figure that is identical to the figure that was published. The documentation of X-PSI describes the postprocessing module of the software. However, the documentation describes v0.5 of X-PSI, which is backward incompatible with v0.1 used in the original analysis. We found that the postprocessing module of X-PSI failed to process the output files from MULTINEST. Instead, we used the postprocessing modules and scripts from PyCBC Inference¹²—a Python toolkit for the Bayesian analysis of gravitational wave signals—to plot all of the posteriors. We converted the .dat files produced by the MULTINEST sampler into PyCBC-readable .hdf files and used the `pycbc_plot_posterior` script on these files. This emphasizes the importance of releasing the set of all of the executables used in the original analysis, including the postprocessing and the plotting scripts.

In addition to reproducing the original analyses, we also explored the effects of using broader prior bounds for the radius. The original analysis put the upper bound at 16 km for the neutron star's radius. We changed it to 25 km and found that the posteriors are unaffected.

TABLE 1. Summary of job statistics for the original analysis published by Riley et al.⁴ in the Zenodo repository and the reanalyses.*

Model		Results in Riley et al. ⁴	Data From Zenodo Release	Reproducibility Analysis	Analysis Using Broader Radius Priors	Analysis Using 4000 Sampler Live Points
ST+PST	$\ln Z$	-36368.28	-36366.65	-36365.52	-36365.24	-36364.39
	Mass ($M_{\odot}$)	$1.34^{+0.15}_{-0.16}$	$1.34^{+0.15}_{-0.15}$	$1.34^{+0.16}_{-0.15}$	$1.36^{+0.16}_{-0.16}$	$1.35^{+0.16}_{-0.16}$
	Equatorial radius (km)	$12.71^{+1.14}_{-1.19}$	$12.7^{+1.1}_{-1.2}$	$12.8^{+1.2}_{-1.2}$	$12.9^{+1.2}_{-1.2}$	$12.9^{+1.3}_{-1.2}$
	Likelihood evaluations	78,343,018	78,343,018	157,814,515	139,593,698	589,513,174
	Nested replacements	57,972	57,972	56,896	56,596	225,856
	Weighted posterior samples	20,177	12,242	11,896	11,749	46,488
	CPU hours	42,453	42,453	48,384	55,296	179,712
	Number of cores	960	960	288	384	288
ST+EST	$\ln Z$	-36,367.81	-36,366.17	-36,366.14	-36,366.16	—
	Mass ($M_{\odot}$)	$1.46^{+0.17}_{-0.18}$	$1.46^{+0.17}_{-0.18}$	$1.46^{+0.17}_{-0.17}$	$1.47^{+0.19}_{-0.19}$	—
	Equatorial radius (km)	$13.89^{+1.14}_{-1.30}$	$13.9^{+1.1}_{-1.3}$	$13.8^{+1.2}_{-1.2}$	$14^{+1.4}_{-1.4}$	—
	Likelihood evaluations	88,965,106	88,965,106	89,850,127	143,920,078	—
	Nested replacements	53,149	53,149	53,098	52,358	—
	Weighted posterior samples	20,177	12,242	10,944	10,828	—
	CPU hours	61,210	61,210	48,384	55,296	—
	Number of cores	960	960	288	480	—

*We show the key results in the first section of the models ST+PST and ST+EST for the Bayesian evidence ($\ln Z$) obtained from the analysis as well as the measurement of the mass and equatorial radius of the target pulsar. The results in Riley et al.⁴ are obtained after postprocessing the data they obtained from the job. The latter has been publicly released in the Zenodo repository, so the computational details (CPU hours and number of cores) for these two columns are the same. In contrast, the evidence values and measurements of mass and radius differ.

This test is useful for the model ST+EST, where the posteriors are cut off at the upper bound of the prior. We aimed to check if the posteriors were affected if the prior base was increased. Since the posteriors did not change significantly, we conclude that the data observed by the NICER instrument is informative and that our choice of priors does not heavily influence the Bayesian analysis. The posteriors for the analyses with broader radius priors for models ST+PST and ST+EST are shown in green in Figures 3 and 4, respectively.

For the model ST+PST, we also perform an analysis with an increased number of live points used by the sampler to sample the posterior probability distribution. The original analysis used 1000 live points, and we increased it to 4000 live points to check the robustness of the result to the sampler configuration. The posteriors do not significantly change when the number of live points for the sampler increases.

LESSONS LEARNED

We compiled a list of challenges we encountered while reproducing the analysis and noted the lessons learned. We discuss the guidelines to make computational analysis, such as that done by Riley et al.,⁴ reproducible. Table 2

lists whether the data, software, and documentation components were available, incomplete, or unavailable before our reproducibility study.

- **Input data availability:** The raw data for the NICER observation of PSR J0030+451 were not made available by Riley et al.⁴ through their Zenodo repository. However, processed data were included in the Zenodo release and had accompanying documentation.
- **Software availability:** Riley et al.⁴ use X-PSI v0.1 to analyze PSR J0030+451 data. The code is open source and publicly available on GitHub.
- **Software documentation:** X-PSI comes with extensive, publicly accessible documentation. Since the framework of the code is modular, the documentation goes over each module in depth, explaining the physics associated with it and providing examples. However, the code and its documentation have evolved significantly since they were used for the original analysis. The documentation during our reproducibility effort relates to X-PSI v0.5, whereas the original analysis used X-PSI v0.1.




FIGURE 3. Comparison of posterior probability distributions for the mass, radius, and compactness of J0030 obtained by Riley et al.⁴ (blue), obtained when reproducing the analysis (orange), and for the analysis with the broader radius priors (green) using the hot spot model ST+PST. The corner plot shows the three parameters' 1- and 2-D marginal posteriors. The priors used for the reanalysis are the same as in the original analysis.

- › *Software installation and dependencies:* The instructions for the installation of X-PSI include information about all of the software dependencies. They also provide .yaml files that can be used to create a virtual environment with the basic dependencies resolved. The installation manual has clear instructions for installing the sampler and the parallelization software.
- › *Configuration files:* The Zenodo repository has all of the configuration files used by X-PSI to generate the hot spot models. The availability of the configuration files was crucial to successfully reproducing the results. The configuration files shared in the Zenodo repository streamline the setup of the jobs. Combined with the documentation, modifying the original analysis and changing the sampler configuration and the prior bounds for the radius was easy.
- › *Computational resources:* The original analysis of Riley et al.⁴ used the Dutch national

SURFsara supercomputer Cartesius. As is common in attempts to reproduce analyses, we did not have access to these computational resources or the original environment used by Riley et al. To execute X-PSI on the large-scale resources available to us, we had to adapt the X-PSI deployment to fit a different scheduler and create an overlay container that could run on this cluster. By demonstrating that this is possible, we show that overcoming the barrier to reproducibility presented by the lack of access to computing resources is possible.

- › *Postprocessing scripts:* The unavailability of the postprocessing scripts to analyze the output data from the analysis and plot the posteriors made it unfeasible for us to generate the same plot as the one present in the publication of the original results. We had to use PyCBC software to plot the posterior distributions. Postprocessing scripts are a crucial part of reproducibility




FIGURE 4. Comparison of posterior probability distributions for the mass, radius, and compactness of J0030 obtained by Riley et al.⁴ (blue), obtained when reproducing the analysis (orange), and for the analysis with the broader radius priors (green) using the hot spot model ST+EST. The corner plot shows the three parameters' 1- and 2-D marginal posteriors. The priors used for the reanalysis are the same as in the original analysis.

in the software workflow. The scripts to process the raw data are also absent.

- **Output data availability:** The Zenodo repository included the posterior output files for all of the analysis jobs performed by the original study's authors. The output files included all the files generated by the sampler MULTINEST, including the posterior file and the history of all the sample points throughout the analysis. It also includes the output of the jobs—including the job scheduler logs and error messages generated by X-PSI during the analysis.

CONCLUSION

Conducting reproducible research is an essential step toward open science. In this article, we described the procedure and challenges involved in reproducing the measurement of the mass and radius of the pulsar PSR J0030 + 451 from the X-ray data observed by NICER.

Given the release of the Zenodo repository containing the data and the configuration scripts used for the original analysis, we were able to reproduce the analysis by Riley et al.⁴ to measure the mass and the radius of PSR J0030 + 451. The postprocessing scripts plot the posteriors using the output file produced by X-PSI, which is absent. We could not use the code and its documentation to plot the posteriors as shown in the original publication. Instead, after converting the output file to an .hdf file, we used the postprocessing module of PyCBC to plot the posteriors. This highlights the importance of releasing the entire set of scripts, from data processing to postprocessing of the analysis output, in a containerized format to reproduce the analysis.

Apart from reproducing the measurement, we changed the prior probabilities of the radius from the original analysis, increasing the upper bound from 16 to 25 km. Despite the broader range of possible radii from the prior, we get the exact posterior distribution as the

TABLE 2. Availability of data, scripts, code, and documentation before our reproducibility study.*

Data		
Raw input data	Unavailable	—
Processed input data	Available	https://doi.org/10.5281/zenodo.5506838
Output data	Available	https://doi.org/10.5281/zenodo.5506838
Software		
Code	v0.1 available ³	https://github.com/xpsi-group
Documentation	Available for v0.5	https://xpsi-group.github.io/xpsi/index.html
Software dependencies	Available	—
Configuration files	Available	—
Postprocessing scripts	Unavailable	—

*v: version.

original analysis. We also increased the number of points used by the sampler from 1000 to 4000 and found no significant change in the posterior probability distribution.

Our work also shows that it is possible to reproduce analyses that require large-scale computational resources without access to the original hardware. This is significant, as access to resources is often a major barrier to reproducibility. Scientists wishing to reproduce findings might not have allocations on the original resources, or the original resource may have been decommissioned. Using the Singularity overlay container shows that executing MPI code across a heterogeneous cluster that uses a different operating system than the original hardware is possible.

To aid future researchers who want to reproduce the analysis of PSR J0030 + 451 data, the Docker image created for our analysis is publicly available.^e The specific tag of the container's image used for the reproducibility analyses is "8d3b23d." The Docker file is also available publicly on the GitHub repository accompanying this article.^f We provide the postprocessing script and the PyCBC installation required to produce the posterior corner plots.

ACKNOWLEDGMENTS

We thank Anna Watts for the helpful discussions. This work was supported by the U.S. National Science Foundation Grant 2207264, Grant 2041977, Grant 2041901, Grant 2028923, Grant 2028930, Grant 1841399, and Grant 1941443. The work of Ingo Tews was supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics Contract DE-AC52-06NA25396 and by the Laboratory Directed Research and Development Program of Los Alamos National Laboratory Project 20220541ECR and Project 20230315er. Syracuse University provided computational resources.

^e<https://hub.docker.com/r/chaitanyaafle/nicer>

^f<https://github.com/sugwg/nicer-reproducibility-J0030>

REFERENCES

1. National Academies of Sciences Engineering and Medicine, *Reproducibility and Replicability in Science*. Washington, DC, USA: National Academy Press, 2019. [Online]. Available: <https://nap.nationalacademies.org/catalog/25303/reproducibility-and-replicability-in-science>
2. K. C. Gendreau et al., "The Neutron star Interior Composition Explorer (NICER): Design and development," in *Proc. SPIE Space Telescopes Instrum., Ultraviolet Gamma Ray*, J.-W. A. den Herder, T. Takahashi, and M. Bautz, Eds., Jul. 2016, vol. 9905, pp. 420–435, doi: [10.1117/12.2231304](https://doi.org/10.1117/12.2231304).
3. T. E. Riley, "X-PSI: X-ray pulse simulation and inference," Astrophysics Source Code Library, record ascl:2102.005, Feb. 2021.
4. T. E. Riley et al., "A NICER view of PSR J0030 + 0451: Millisecond pulsar parameter estimation," *Astrophys. J. Lett.*, vol. 887, no. 1, 2019, Art. no. L21, doi: [10.3847/2041-8213/ab481c](https://doi.org/10.3847/2041-8213/ab481c).
5. M. C. Miller et al., "PSR J0030 + 0451 mass and radius from NICER data and implications for the properties of neutron star matter," *Astrophys. J. Lett.*, vol. 887, no. 1, 2019, Art. no. L24, doi: [10.3847/2041-8213/ab50c5](https://doi.org/10.3847/2041-8213/ab50c5).
6. D. A. Brown, K. Vahi, M. Taufer, V. Welch, and E. Deelman, "Reproducing GW150914: The first observation of gravitational waves from a binary black hole merger," *IEEE Comput. Sci. Eng.*, vol. 23, no. 2, pp. 73–82, Mar./Apr. 2021, doi: [10.1109/MCSE.2021.3059232](https://doi.org/10.1109/MCSE.2021.3059232).
7. R. Patel et al., "Reproducibility of the first image of a black hole in the galaxy M87 from the event horizon telescope collaboration," *IEEE Comput. Sci. Eng.*, vol. 24, no. 5, pp. 42–52, Sep./Oct. 2022, doi: [10.1109/MCSE.2023.3241105](https://doi.org/10.1109/MCSE.2023.3241105).
8. R. Ketron et al., "A case study in scientific reproducibility from the event horizon telescope (EHT)," in *Proc. 20th IEEE Int. Conf. eScience*,

- Innsbruck, Austria: IEEE Comput. Soc. Press, Sep. 2021, pp. 249–250, doi: [10.1109/eScience51609.2021.00045](https://doi.org/10.1109/eScience51609.2021.00045).
9. F. Feroz, M. P. Hobson, and M. Bridges, “MultiNest: An efficient and robust Bayesian inference tool for cosmology and particle physics,” *Monthly Notices Roy. Astronomical Soc.*, vol. 398, no. 4, pp. 1601–1614, Oct. 2009, doi: [10.1111/j.1365-2966.2009.14548.x](https://doi.org/10.1111/j.1365-2966.2009.14548.x).
 10. L. Dalcín, R. Paz, and M. Storti, “MPI for Python,” *J. Parallel Distrib. Comput.*, vol. 65, no. 9, pp. 1108–1115, 2005, doi: [10.1016/j.jpdc.2005.03.010](https://doi.org/10.1016/j.jpdc.2005.03.010). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0743731505000560>
 11. E. Gabriel et al., “Open MPI: Goals, concept, and design of a next generation MPI implementation,” in *Recent Advances in Parallel Virtual Machine and Message Passing Interface*. Berlin, Germany: Springer-Verlag, 2004, pp. 97–104.
 12. C. M. Biwer et al., “PyCBC inference: A Python-based parameter estimation toolkit for compact binary coalescence signals,” *Publications Astronomical Soc. Pacific*, vol. 131, no. 996, 2019, Art. no. 024503, doi: [10.1088/1538-3873/aaf0b](https://doi.org/10.1088/1538-3873/aaf0b).

CHAITANYA AFLE is a Ph.D. candidate in physics at Syracuse University, Syracuse, NY, 13244, USA. His research interests include gravitational wave astronomy and astrophysics, data analysis, and multimessenger astronomy. Contact him at chafle@syr.edu.

PATRICK R. MILES is with Syracuse University, Syracuse, NY, 13244, USA. Miles received his M.S. degree in physics from Syracuse University in 2021. Contact him at pmiles@g.syr.edu.

SILVINA CAÍNO-LORES is a research assistant professor at the University of Tennessee, Knoxville, TN, 37996, USA. Her research interests include cloud computing, in-memory computing and storage, high-performance computing scientific simulations, and datacentric paradigms. Caíno-Lores received her Ph.D. degree in computer science and technology from Carlos III University of Madrid, Spain. She is a Member of IEEE. Contact her at scainolo@utk.edu.

COLLIN D. CAPANO is a senior scientist at the Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476, Potsdam, Germany, and a high-performance computing facilitator at the University of Massachusetts Dartmouth, Dartmouth, MA, 02747, USA. His research interests include the analysis of gravitational waves from black holes and

neutron stars. Capano received his Ph.D. degree in physics from Syracuse University. Contact him at collin.capano@aei.mpg.de.

INGO TEWS is a staff scientist at Los Alamos National Laboratory, Los Alamos, NM, 87545, USA. His research interests include fundamental interactions between neutrons and protons in atomic nuclei and neutron stars. Tews received his Ph.D. degree in physics from Technische Universität Darmstadt. Contact him at itews@lanl.gov.

KARAN VAHI is a senior computer scientist at the Information Sciences Institute, University of Southern California, Los Angeles, CA, 90007, USA. His research interests include scientific workflows and distributed computing systems. Vahi received his M.S. degree in computer science from the University of Southern California. Contact him at vahi@isi.edu.

EWA DEELMAN is a research director at University of Southern California/Information Sciences Institute; is a research professor with the Computer Science Department, University of Southern California, Los Angeles, CA, 90007, USA; and leads the design and development of the Pegasus Workflow Management software. Her research interests include the interplay between automation and the management of scientific workflows, including resource provisioning and data management. Deelman received her Ph.D. degree in computer science from Rensselaer Polytechnic Institute. She is a Fellow of IEEE. Contact her at deelman@isi.edu.

MICHELA TAUFER holds the Dongarra Professorship in High-Performance Computing in the Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, TN, 37996, USA. Her research interests include the intersection of computational sciences, high-performance computing, and data analytics. Tauber received her Ph.D. degree in computer science from the Swiss Federal Institute of Technology. She is a Senior Member of IEEE. Contact her at tauber@utk.edu.

DUNCAN A. BROWN is the Charles Brightman Professor of Physics at Syracuse University, Syracuse, NY, 13244, USA. His research interests include gravitational wave astronomy, astrophysics, and large-scale scientific workflows. Brown received his Ph.D. degree in physics from the University of Wisconsin-Milwaukee. He is a fellow of the American Physical Society. Contact him at dabrown@syr.edu.