

Bayesian approach to radio frequency interference mitigationS. A. K. Leeney¹,^{*,§} W. J. Handley¹,^{†,§} and E. de Lera Acedo¹,^{‡,§}*The University of Cambridge, Astrophysics Group, Cavendish Laboratory,
J. J. Thomson Avenue, Cambridge, CB3 0HE, United Kingdom* (Received 5 May 2023; accepted 29 August 2023; published 29 September 2023)

Interfering signals such as radio frequency interference from ubiquitous satellite constellations are becoming an endemic problem in fields involving physical observations of the electromagnetic spectrum. To address this we propose a novel data cleaning methodology. Contamination is simultaneously flagged and managed at the likelihood level. It is modeled in a Bayesian fashion through a piecewise likelihood that is constrained by a Bernoulli prior distribution. The techniques described in this paper can be implemented with just a few lines of code.

DOI: [10.1103/PhysRevD.108.062006](https://doi.org/10.1103/PhysRevD.108.062006)**I. INTRODUCTION**

Satellite constellations in low earth orbit such as SpaceX's Starlink will likely number 1,00,000 by 2030 [1]. Described in a recent Nature article as “horrifying” [2], the impact of radio frequency interference (RFI) created by these “megaconstellations” on astronomy is of significant concern. Interfering signals like RFI are problematic because they cause information in contaminated frequency channels to be lost which can lead to significant systematic error if not properly modeled.

The frequency bands that these satellites and other modern telecommunications devices operate in are not protected or reserved for astronomy. They clash with current and next generation experiments such as the ngVLA [3] and the SKA [4], which plan probe yet unseen epochs of the Universe. Furthermore, the latest telescopes operate in large regions of the sky and across wide frequency bands, making it impossible to avoid these satellite swarms. This is pushing astronomy to remote corners of earth and beyond, to space [5].

However, moving projects to space is not a long term solution. There is only a short time window until the RFI-quiet dark side of the moon becomes contaminated by satellites and other devices required by projects such as

LuSEE [6] and the LCRT [7] which themselves are hoping to exploit the clear lunar skies. Furthermore, space itself is filled with cosmic rays, which lead to RFI-like interference and thus have been problematic for the James Web Space Telescope's near infrared spectrometer [6].

Consequentially, RFI is becoming the fundamental bottleneck in modern astronomy and beyond; with modern observations being at the forefront of many fields of fundamental physics, such as dark matter detection. New statistical techniques are therefore urgently required to address this rapidly growing problem.

Furthermore, projects such as the Event Horizon Telescope [8], Planck [9], and BICEP [10] all used Bayesian statistics in their data analysis pipelines, but there is currently no way to manage RFI in a Bayesian sense forcing astronomers to patch traditional RFI mitigation algorithms into their Bayesian systems. Modern projects implement singular value decomposition [11] watershed segmentation [12], deep learning methods [13–15], and more recently Gibbs sampling [16].

Scientists are constantly searching for more elusive and faint signals, thus data cleaning is becoming increasingly important. The Laser Interferometer Gravitational-wave Observatory (LIGO) [17], for example, is sensitive enough to detect a car starting miles away and thus is extremely sensitive to data corruption [18]. The next generation aLIGO [19] will be even more so. As such there is a need for novel data cleaning methodologies beyond fields involving measurements of the electromagnetic spectrum.

In this paper we propose a data cleaning methodology that takes a Bayesian approach, where contamination is both flagged and managed at the likelihood level. Detailed benchmarking of this technique will be provided when it is applied on data from a next generation low frequency global experiment, REACH [20], which will see first light

*saki2@cam.ac.uk

†wh260@cam.ac.uk

‡ed330@cam.ac.uk

§Also at Kavli Institute for Cosmology, Madingley Road, Cambridge, CB3 0HA, United Kingdom.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

this year. A usage example of the methodologies described in this work can be found in [21].

II. THEORY

A. Bayesian inference

Bayesian methods can be used to perform parameter estimates and model comparison. A model \mathcal{M} uses data \mathcal{D} to infer its free parameters θ . Using Bayes theorem,

$$P(\mathcal{D}|\theta) \times P(\theta) = P(\theta|\mathcal{D}) \times P(\mathcal{D}), \quad (1)$$

$$\mathcal{L} \times \pi = \mathcal{P} \times \mathcal{Z}, \quad (2)$$

the prior π is updated onto the posterior \mathcal{P} in light of the likelihood \mathcal{L} and furthermore the Bayesian evidence \mathcal{Z} can be inferred by computing the integral

$$\mathcal{Z} = \int \mathcal{L}(\theta) \times \pi(\theta) d\theta. \quad (3)$$

In practice, \mathcal{P} and \mathcal{Z} can be determined simultaneously using a Bayesian numerical solver. We use the nested sampling algorithm POLYCHORD [22], where a series of points generated within π are updated such that they sequentially contract around the peak(s) of the likelihood, forming the posterior which can be used to generate parameter estimates. The artefacts of this process can then be used to compute \mathcal{Z} , which is used for model comparison. For a more detailed description of Bayesian inference and nested sampling see [23].

B. Data cleaning likelihood

A sufficiently contaminated data point can be considered corrupted. Any information relevant to the model is lost and furthermore it cannot be modeled as Gaussian noise. Assuming \mathcal{D} is uncorrelated, the likelihood

$$\mathcal{L} = P(\mathcal{D}|\theta) = \prod_i \mathcal{L}_i(\theta), = \prod_i P(\mathcal{D}_i|\theta), \quad (4)$$

where i represents the i th data point, is insufficient to model such contaminated data. It is therefore necessary to model the likelihood that each data point is corrupted.

Thus, we introduce a piecewise likelihood including the possibility of corruption of data:

$$P(\mathcal{D}_i|\theta) = \begin{cases} \mathcal{L}_i(\theta) & \text{:uncorrupted} \\ \Delta^{-1}[0 < \mathcal{D}_i < \Delta] & \text{:corrupted.} \end{cases} \quad (5)$$

Corruption is modeled as the data becoming completely unreliable and therefore being distributed uniformly within some range Δ (which, as a scale of corruption, has the same dimensions as the data). An efficient way to write this likelihood is

$$P(\mathcal{D}|\theta, \varepsilon) = \prod_i \mathcal{L}_i^{\varepsilon_i} \Delta^{\varepsilon_i-1}, \quad (6)$$

where the Boolean mask vector ε has a i th component which takes the value 1 if the datum i is uncorrupted and value 0 if corrupted.

We do not know before the data arrive whether or not they are corrupted. We may infer this in a Bayesian fashion, by ascribing a Bernoulli prior probability p_i of corruption (which has dimensions of probability); i.e.,

$$P(\varepsilon_i) = p_i^{(1-\varepsilon_i)} (1 - p_i)^{\varepsilon_i}. \quad (7)$$

Both Δ and p_i are required for a dimensionally consistent analysis. It should be noted that above we assume the *a priori* probability that each bin is corrupted is uncorrelated, i.e., $P(\varepsilon) = \prod_i P(\varepsilon_i)$, which in practice will almost certainly not be true. We will discuss later the extent to which this assumption can be considered valid.

Multiplying Eqs. (6) and (7) yields

$$P(\mathcal{D}, \varepsilon|\theta) = \prod_i [\mathcal{L}_i(1 - p_i)]^{\varepsilon_i} [p_i/\Delta]^{(1-\varepsilon_i)} \quad (8)$$

and to recover a likelihood independent of ε we formally can marginalize out:

$$\begin{aligned} P(\mathcal{D}|\theta) &= \sum_{\varepsilon \in \{0,1\}^N} P(\mathcal{D}, \varepsilon|\theta) \quad (9) \\ &= \sum_{\varepsilon \in \{0,1\}^N} \prod_i [\mathcal{L}_i(1 - p_i)]^{\varepsilon_i} [p_i/\Delta_i]^{(1-\varepsilon_i)}. \quad (10) \end{aligned}$$

This would require the computation of the all 2^N terms in Eq. (10). For realistic values of N , this computation becomes impractical. However, if it is assumed that the most likely model [i.e., the maximum term in Eq. (10)] dominates over the next to leading order terms, we can make the approximation

$$P(\mathcal{D}, \varepsilon|\theta) \approx \delta_{\varepsilon, \varepsilon^{\max}} \times P(\mathcal{D}, \varepsilon^{\max}|\theta), \quad (11)$$

where δ_{ij} is the usual Kronecker delta function, and ε^{\max} is the mask vector which maximizes the likelihood $P(\mathcal{D}, \varepsilon|\theta)$, namely,

$$\varepsilon_i^{\max} = \begin{cases} 1, & \mathcal{L}_i(1 - p_i) > p_i/\Delta_i \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

Under this approximation we find that the sum in Eq. (10) becomes

$$P(\mathcal{D}|\theta) \approx P(\mathcal{D}, \varepsilon^{\max}|\theta). \quad (13)$$

In practice the approximation in Eq. (13) is only valid if the next to leading order term is much smaller, such that

$$P(\mathcal{D}|\theta, \varepsilon^{\max}) \gg \max_j P(\mathcal{D}|\theta, \varepsilon^{(j)}), \quad (14)$$

where $\varepsilon^{(j)}$ is ε^{\max} with its j th bit flipped:

$$\varepsilon_k^{(j)} = \begin{cases} 1 - \varepsilon_k^{\max} & k = j \\ \varepsilon_k^{\max} & k \neq j \end{cases} \quad (15)$$

and we may use Eq. (14) as a consistency check.

To summarize, we can correct for contamination under these approximations by replacing the original likelihood $\mathcal{L} = \prod_i \mathcal{L}_i$ in Eq. (4) with

$$P(\mathcal{D}|\theta) = \prod_i [\mathcal{L}_i(1 - p_i)]^{\varepsilon_i^{\max}} [p_i/\Delta]^{(1-\varepsilon_i^{\max})}, \quad (16)$$

where ε^{\max} is defined by Eq. (12).

$$\log P(\mathcal{D}|\theta) = \begin{cases} \log \mathcal{L}_i + \log(1 - p_i), \\ \log p_i - \log \Delta, \end{cases}$$

The corrected likelihood is then updated iteratively via the selected Bayesian sampling method, compressing the prior onto the posterior while simultaneously correcting for contamination. One may also notice that the condition $\log \mathcal{L}_i + \log(1 - p_i) > \log p_i - \log \Delta$ in Eq. (18) relates to a Logit function, such that

$$\log \mathcal{L}_i + \log \Delta > \text{LOGIT}(P). \quad (19)$$

Logit functions are used routinely as an activation function in binary classification tasks, hinting at the potential of a future extension of this work using machine learning.

III. TOY EXAMPLE

We will initially test this approach on a simple toy model with a basic contaminant signal injected. We then move on to a more realistic and complex case in Sec. IV.

A. Initial testing

Two simple datasets are generated for comparison consisting of a line with $m = 1$ and $c = 1$ with $\sigma = 5$ order Gaussian noise. One is contaminated by an RFI like signal and the other (the ground truth) is not. They are fit in a Bayesian sense, attempting to recover the two free

C. Computing the posterior

The posterior and evidence are computed from Eq. (16) via nested sampling (although any numerical Bayesian sampling method could be used). Taking logs for convenience gives

$$\log P(\mathcal{D}|\theta) = \sum_i [\log \mathcal{L}_i + \log(1 - p_i)] \varepsilon_i^{\max} + [\log p_i - \log \Delta](1 - \varepsilon_i^{\max}), \quad (17)$$

yielding a masked chi squared like term which can be used to distinguish whether there is a statistically significant difference between the classes of data, i.e., corrupted or noncorrupted. Furthermore, the second term in Eq. (17) introduces an Occam penalty. Each time a data point is predicted to be contaminated its likelihood is replaced with the penalty rather than being completely removed. Without this term, the likelihood where all data points are flagged would be larger and thus “more likely” than all other possibilities. Therefore, flagging all datum would always be preferable.

We compute this by imposing the condition in Eq. (12) on Eq. (17) as follows:

$$\begin{cases} \log \mathcal{L}_i + \log(1 - p_i) > \log p_i - \log \Delta \\ \text{otherwise.} \end{cases} \quad (18)$$

parameters m and c using the correcting likelihood in Eq. (18) with

$$\mathcal{L}_i = -\frac{\log(2\pi\sigma^2)}{2} - \frac{[y_i - y_S(x_i; m, c)]^2}{2\sigma^2}, \quad (20)$$

where $\theta = m, c, \sigma$, y_i is the simulated data and $y_S(x_i; m, c)$ are the parameter estimates at the i th sampling iteration used to compute the model $y_i = mx_i + c$. We set $\Delta = \mathcal{D}_{\text{MAX}}$, to encapsulate the full range of possible data values. Δ could likely be fit as a free parameter, as will be investigated further in future works.

Evaluating the posterior on ε we can assess how frequently across the entire sampling run each datum was believed to fit (noncorrupted) or not fit (corrupted) the model. Contaminated points should make up a near zero fraction of the posterior. Conversely, points that are not contaminated would often fit the model and as such contribute significantly to the final posterior distribution. There can also be some points that lie somewhere in between, which the model is less confident are uncontaminated. This may occur with datum that deviate the most from the true signal due to higher order Gaussian noise, for example.

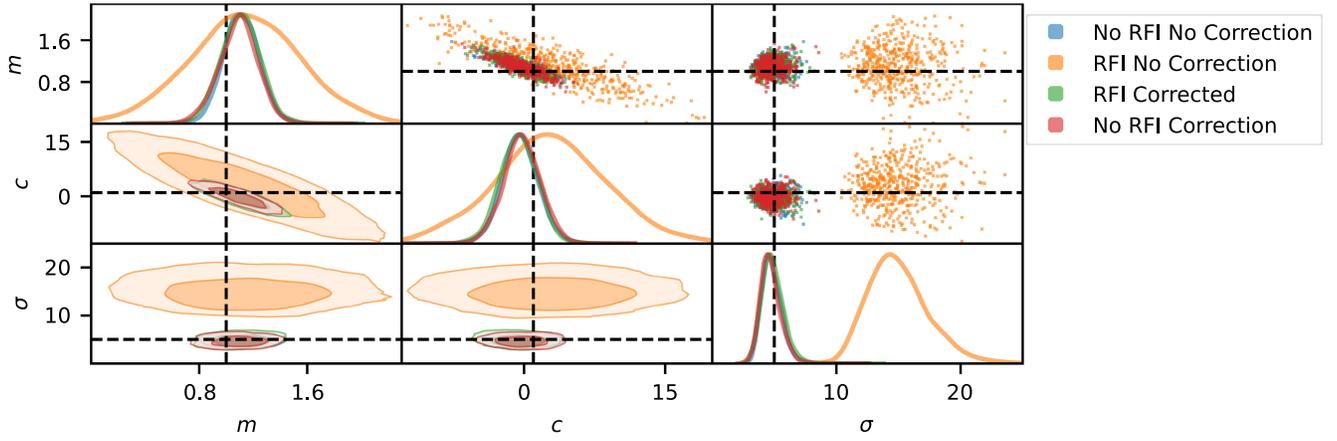


FIG. 1. Showing the parameter distributions inferred from the dataset described in Sec. III A. The top left to bottom right panes show probability distribution functions for m , c , and σ , respectively. Plots are generated using posterior plotting tool ANESTHETIC [24].

It should be emphasized that although ε_i is constrained to binary values, the subsequent mask on ε is not. Unlike traditional RFI flagging algorithms, points are not classified in a binary manner. The mask takes the weighted mean across the posterior. Thus, points more likely to contain RFI will have less “impact” on the final posterior distribution than points believed to be uncontaminated. The mask could be thought of as being slightly opaque to these data points, accounting for the models uncertainty. Incorporating the models confidence in its correction directly into the subsequent parameter estimates makes this approach unique in comparison with its counterparts.

B. Model evaluation

The two aforementioned datasets are evaluated when fit using the likelihood capable of correcting for contamination and also when using a standard likelihood, which cannot natively account for contamination. This generates a total of four posterior distributions for comparison. Of these, all but the contaminated, uncorrected case would be expected to perform similarly if RFI has been effectively mitigated. From a Bayesian standpoint, the simplest model will always be preferable. Thus, for the clean dataset it would be expected that the standard likelihood would be preferred slightly over the correcting likelihood. Figure 1 shows the parameter distributions inferred from the data. The results are as expected; parameters in all but the “RFI No correction” are inferred to within 1σ of their true value.

As seen in Fig. 2, the model that does not correct for RFI is slightly preferred for uncontaminated data. Conversely the correcting model is strongly preferred on the contaminated data indicating that the correction is working as predicted.

Viewing the posterior plots of $P(y|x, \mathcal{D})$ in Fig. 2 it is clear that when RFI is not corrected, the true parameter values are outside the 1σ and sometimes 2σ confidence bounds. Conversely the other three cases fit almost entirely

within the 1σ bounds, indicating that the RFI has been mitigated.

C. Evaluating the $\log p$ dependence

Proper selection of the probability thresholding term $\log p$ is essential. From a Bayesian standpoint it should be set to represent our prior degree of belief in there being RFI in each datum. We assess the $\log p$ dependence while varying $\log p$ as a function of the root-mean-square error (RMSE) on the fit generated from the parameter estimates, the log Bayesian evidence and the mean number of points flagged across all samples.

For high $\log p$, the RMSE is high and we observe in Fig. 3 that the model generates less accurate parameter estimates. Here the threshold is so high that the model is more confident that any of the points are RFI than non-RFI. This matches the corresponding low evidence. The RMSE

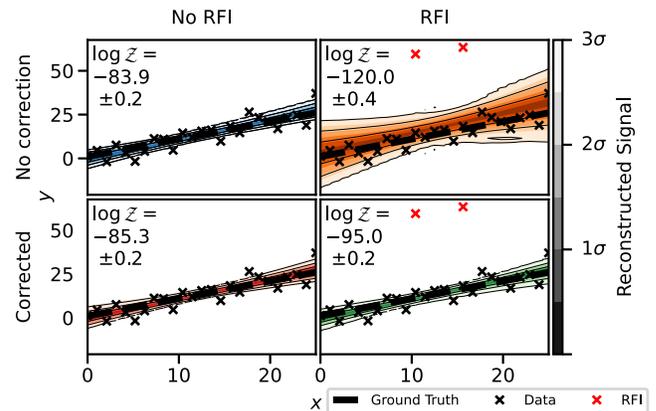


FIG. 2. Showing the inferred parameter estimates in a contour plot, where darker tones indicate higher σ confidence in the parameter estimates. Generated from the dataset described in Sec. III A. The Bayes factor is -1.4 ± 0.3 for the no RFI case and 25.0 ± 0.4 for the RFI case. The plots are generated using the functional posterior plotter FGIVENX [25].

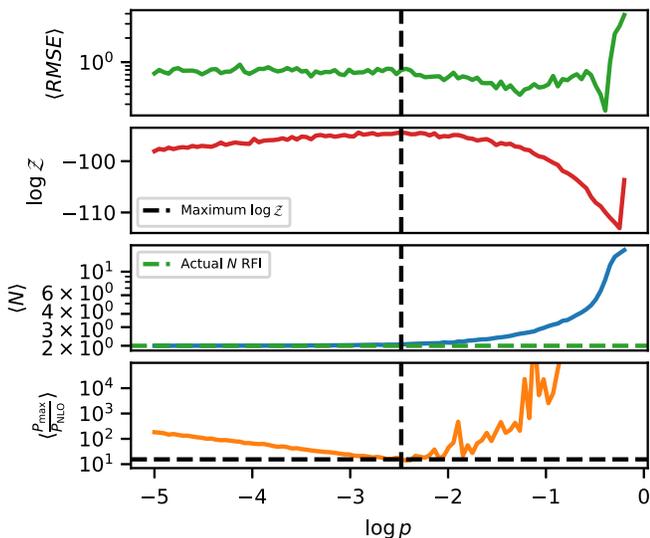


FIG. 3. Assessing how various methods of model evaluation vary as a function of $\log p$. From top to bottom: the RMSE, the log of the Bayesian evidence, then the weighted average number of points flagged and finally the ratio between P_{MAX} and P_{NLO} . All dependant variables (excluding $\log \mathcal{Z}$) are averaged over the weighted posterior samples. The noise observable in these plots is sampling noise; the noise in the simulated data is seeded.

drops as $\log p$ decreases to near its minimum. The model incorrectly flags ≈ 5 data points while the RMSE is low, showing that the model is able to generate accurate parameter estimates whilst over flagging, indicating it is insensitive to false positives. As $\log p$ decreases further, the model is better able to distinguish between higher order Gaussian noise and as such the average number of points predicted to be RFI approaches the true value. As this happens the evidence also reaches its maximum, which indicates that the Bayesian evidence is appropriately showing how well each of the many models created by different $\log p$ values fit the data.

D. To what extent is $P(\mathcal{D}|\theta) \approx P(\mathcal{D}|\theta, \epsilon_{\text{max}})$ valid?

A key assumption is made in Eq. (13) that the leading order term, Eq. (17), is considerably larger than all the other possible terms for $\epsilon \in (0, 1)^N$. It is necessary to test the validity of this approximation by computing Eq. (18) and comparing the result (P_{MAX}) with the next leading order term (P_{NLO}) as calculated by Eq. (14). For $-5 < \log p < -0.1$. These results are displayed in the bottom panel of Fig. 3. P_{MAX} is 18 times larger than P_{NLO} at peak $\log \mathcal{Z}$ and increases linearly for $\log p$ below this. Depending on the $\log p$ selection strategy, P_{MAX} is at least 11 times more likely than the next leading order term. Assuming an appropriate $\log p$ selection strategy, the ratio would be higher, indicating that $P(\mathcal{D}|\theta) \approx P(\mathcal{D}|\theta, \epsilon_{\text{max}})$ is valid.

E. Selection strategy for $\log p$

Various selection strategies could be taken to select the optimal $\log p$ value. For each $\log p$, the model changes. As such, selecting the $\log p$ that maximizes the evidence seems to be the most obvious selection strategy. In the case of the toy model, the peak $\log \mathcal{Z}$ occurs where $\log p = -2.7$. Here, P_{MAX} is 18 times larger than P_{NLO} . Another possible strategy could be to select $\log p$ where the number of points flagged is at its minimum.

It is also possible to ascribe a prior to $\log p$, fitting it as a free parameter thus fully automating the approach. This will be examined further in future works.

IV. REACH EXAMPLE

Finally, we examine a real use case for this method. The REACH [20] radio telescope is designed to detect the 21 cm signal from the cosmic dawn. We select REACH as a testing ground for our methodology because it operates in the same low frequency ranges as the next generation of astronomical experiments, such as the nvGLA or the SKA and takes a Bayesian approach to data analysis [26]. The 21 cm signal is expected to take the shape of an inverted Gaussian, so the model takes the form

$$f(x) = A \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right) \quad (21)$$

with center frequency μ , standard deviation σ and magnitude A all free parameters. The four cases discussed in Sec. III are then examined, but this time on a simulated sky dataset containing a 21 cm signal with two RFI spikes injected. The no RFI correction and ground truth cases are very similar with the simpler (ground truth) case marginally preferred as expected. The RFI corrected case is again similar with a slightly lower evidence due to the penalties

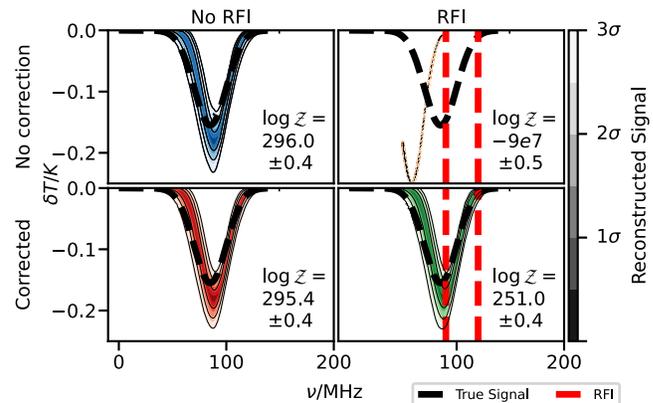


FIG. 4. Showing the results when the RFI correction is applied on simulated data in the REACH data analysis pipeline. The Bayes factor for the no RFI case is -0.6 ± 0.6 and the Bayes factor for the RFI case is $9e7 \pm 0.6$.

incurred during the corrections. The above is also evident when viewing Fig. 4. The reconstructed signal is within 2σ of the true signal for all but the RFI no correction case.

V. CONCLUSIONS

In this paper, which serves as a proof of concept, we show that contamination can be both cleaned and corrected in a fully Bayesian sense at the likelihood level.

We demonstrate our general approach in the context of signal processing for astronomy, but these methods will likely be useful beyond. Forthcoming results from current state of the art low frequency astronomy experiments [27] (where these methods will be benchmarked in the coming months) and extensive testing of this approach as a general Bayesian data cleaning methodology will be used to provide a more detailed analysis in the future.

-
- [1] A. Venkatesan, J. Lowenthal, P. Prem, and M. Vidaurri, *Nat. Astron.* **4**, 1043 (2020).
 - [2] A. Witze, *How Satellite Swarms Pose a Rising Threat to Astronomy* (Nature, Berlin, 2022).
 - [3] M. McKinnon, A. Beasley, E. Murphy, R. Selina, R. Farnsworth, and A. Walter, *Bull. Am. Astron. Soc.* **51**, 81 (2019).
 - [4] I. Prandoni and N. Seymour, [arXiv:1412.6942](https://arxiv.org/abs/1412.6942).
 - [5] P. A. Sabelhaus and J. E. Decker, *Optical, Infrared, and Millimeter Space Telescopes* (SPIE, 2004), Vol. 5487, p. 550.
 - [6] G. Giardino, S. Birkmann, M. Robberto, P. Ferruit, B. J. Rauscher, M. Sirianni, C. A. de Oliveira, T. Boeker, N. Luetzendorf, M. te Plate *et al.*, *Publ. Astron. Soc. Pac.* **131**, 094503 (2019).
 - [7] S. Bandyopadhyay, P. McGarey, A. Goel, R. Rafizadeh, M. Delapierre, M. Arya, N. Chahat, J. Lazio, P. Goldsmith, A. Stoica *et al.*, *lunar Crater Radio Telescope (LCRT) on the Far-Side of the Moon* (Jet Propulsion Laboratory, National Aeronautics and Space, Pasadena, CA, 2021).
 - [8] E. H. T. Collaboration *et al.*, *Astrophys. J. Lett.* **875**, L6 (2019).
 - [9] P. A. Ade, N. Aghanim, M. Alves, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, H. Aussel, C. Baccigalupi *et al.*, *Astron. Astrophys.* **571**, A1 (2014).
 - [10] P. Ade, Z. Ahmed, M. Amiri, D. Barkats, R. B. Thakur, C. Bischoff, J. Bock, H. Boenish, E. Bullock, V. Buza *et al.*, *Phys. Rev. D* **103**, 042002 (2021).
 - [11] A. Offringa, A. De Bruyn, M. Biehl, S. Zaroubi, G. Bernardi, and V. Pandey, *Mon. Not. R. Astron. Soc.* **405**, 155 (2010).
 - [12] J. Kerrigan, P. L. Plante, S. Kohn, J. C. Pober, J. Aguirre, Z. Abdurashidova, P. Alexander, Z. S. Ali, Y. Balfour, A. P. Beardsley *et al.*, *Mon. Not. R. Astron. Soc.* **488**, 2605 (2019).
 - [13] J. Akeret, C. Chang, A. Lucchi, and A. Refregier, *Astron. Comput.* **18**, 35 (2017).
 - [14] A. Vafaei Sadr, B. A. Bassett, N. Oozeer, Y. Fantaye, and C. Finlay, *Mon. Not. R. Astron. Soc.* **499**, 379 (2020).
 - [15] H. Sun, H. Deng, F. Wang, Y. Mei, T. Xu, O. Smirnov, L. Deng, and S. Wei, *Mon. Not. R. Astron. Soc.* **512**, 2025 (2022).
 - [16] F. Kennedy, P. Bull, M. Wilensky, and S. Choudhuri, *Astrophys. J. Suppl. Ser.* **266**, 23 (2023).
 - [17] B. P. Abbott *et al.*, *Rep. Prog. Phys.* **72**, 076901 (2009).
 - [18] R. Ormiston, T. Nguyen, M. Coughlin, R. X. Adhikari, and E. Katsavounidis, *Phys. Rev. Res.* **2**, 033066 (2020).
 - [19] G. M. Harry (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084006 (2010).
 - [20] E. de Lera Acedo, D. de Villiers, N. Razavi-Ghods, W. Handley, A. Fialkov, A. Magro, D. Anstey, H. Bevins, R. Chiello, J. Cumner *et al.*, *Nat. Astron.* **6**, 984 (2022).
 - [21] <https://github.com/samleeny/publications>.
 - [22] W. Handley, M. Hobson, and A. Lasenby, *Mon. Not. R. Astron. Soc.* **453**, 4384 (2015).
 - [23] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial* (Oxford University Press, Oxford, 2006).
 - [24] W. Handley, *J. Open Source Software* **4**, 1414 (2019).
 - [25] W. Handley, *J. Open Source Software* **3**, 849 (2018).
 - [26] D. Anstey, E. de Lera Acedo, and W. Handley, *Mon. Not. R. Astron. Soc.* **506**, 2041 (2021).
 - [27] E. de Lera Acedo, in *2019 International Conference on Electromagnetics in Advanced Applications (ICEAA)* (IEEE, New York, 2019), pp. 0626–0629.