Draft version April 6, 2026
Typeset using LaTeX **twocolumn** style in AASTeX7.0.1

Power Spectrum Injections for the CHIME All-sky Multiday Pulsar Stacking Search

Reynier Squillace [1,2]

[1] *University of Virginia*
[2] *National Radio Astronomy Observatory*

ABSTRACT

The injection of artificial signals is a key tool in time-domain radio astronomy, with particular import to the study of pulsars and Fast Radio Bursts (FRBs). Both hardware and software injections into time series are common, and can help probe observational biases impacting the detection of radio transients. In the case of the CHIME All-sky Multiday Pulsar Stack Search (CHAMPSS), the computational demands of both hardware and software injections into the time series force an alternative approach. We present methods for injecting artificial pulse signals into a decoherent power spectrum. The injection process detailed here is computationally negligible compared to the overhead of the CHAMPSS search pipeline, and can be broadly applied to any power spectrum search.

We applied our injection methods to probing the in-system biases of CHIME and the CHAMPSS search pipeline. We find that CHAMPSS is most sensitive to pulsars with frequencies on the order of $10^0 - 10^2$ Hz, with a single day flux sensitivity down to ~ 0.1 mJy. Our injections show weak dependence on other factors such as dispersion measure (DM), duty cycle, and sky coordinates. These findings suggest that frequency-dependent noise and sampling limitations are the dominant extrinsic factors decreasing CHAMPSS' sensitivity to pulsars.

1. INTRODUCTION

Any survey of radio transients whose goals include population-level science must contend with the sensitivity and completeness of its own search pipeline. We define sensitivity as the probability $p(C|\mathbf{B})$ that a search pipeline will identify as a candidate a signal with some set of characteristics $\mathbf{B}$. The completeness, meanwhile, is the probability $p(C|\mathbf{B}; \sigma \geq \sigma_0)$ that a candidate which *could* be detected given a significance greater than the candidate cutoff significance σ_0 actually gets detected.

The influence of each characteristic in $\mathbf{B}$ on $p(C|\mathbf{B})$ biases any derived population distributions or statistics and the underlying physics these may imply. Instead of representing the true distribution $D(\mathbf{B})$, the survey results will represent a biased distribution $D(\mathbf{B}) \times p(C|\mathbf{B})$. Injecting artificial signals prior to performing a search can compensate for this bias.

The CHIME All-sky Multiday Pulsar Stacking Search (CHAMPSS) searches many days' observations of the radio sky in pursuit of new, faint pulsars. CHAMPSS uses the CHIME/FRB backend, which has already found more than 80 pulsars (D. C. Good et al. 2021; F. A. Dong et al. 2023). However, where CHIME/FRB searches for single pulses, CHAMPSS employs a periodicity search across many days of integrated power spectra. While we lose phase information, searching stacked power spectra allows us to search the entire sky simultaneously for faint pulsars that may not have been found in earlier, more shallow surveys. We previously detailed the CHAMPSS search pipeline and preliminary results in C. Andrade et al. (2025). At the time of this publication, CHAMPSS has found 33 new pulsars, with $\gtrsim 100$ future discoveries expected. We require comprehensive injections of simulated pulsars to draw any population-level conclusions from this predicted pulsar dataset.

The injection of simulated signals into pulsar surveys is not new. P. Lazarus et al. (2015) injected simulated time domain signals into data from the Mock Spectrometer on Arecibo to constrain the selection bias of the PALFA survey. They found a dropoff in sensitivity for frequencies $f \leq 10$ Hz, which they attributed to red-noise. Most previous injection schema, however, were for the study of Fast Radio Bursts (FRBs). Given the open questions surrounding the origin of FRBs, and the potential for population statistics to hint at answers, injections are a core part of most major FRB search programs. CHIME/FRB ran a suite of 85,000 simulated FRB injections into a copy of their timestream data (M. Merryfield et al. 2023), a route closed to us due to unrealistic storage requirements. They found

Email: squillace@virginia.edu

that the chance of detection depended most strongly on fluence and pulse width, with wider pulses less likely to be retrieved. They did not find a significant DM dependence. J.-W. Mao et al. (2022) found similar results for their survey on the Nanshan 26m Radio Telescope, albeit with a bias against DM $\lesssim 400\,\mathrm{pc\,cm^{-3}}$ in one of the three search pipelines they tested.

2. INJECTION ALGORITHMS

In my report last year for VSGC, I summarized the injection pipeline. I will therefore skip it in this report and proceed to results.

3. INJECTION SCHEMA

We performed $\sim$ 80,000 injections into a synthetic power spectrum arising from Gaussian noise timestream data with no rednoise or RFI, $\sim 300,000$ injections total spread across 7 days of single-day power spectra sampled across 3 weeks, and $\sim$ 200,000 injections into 21-day stacks. These data were taken in November of 2026 during an early CHAMPSS commissioning run.

The injection distributions $p(\mathrm{B}_i)$ are uncorrelated such that $p(\mathbf{B}) = p(\mathrm{B}_i) \times p(\mathrm{B}_j) ... \times p(\mathrm{B}_{N-1})$. We drew our injection parameters from the same uncorrelated distributions each time, with the exception of different lower and upper limits on our uniform flux and DM distributions. The lower limit of our flux distribution was 0.01 mJy for single-day injections and 0.001 mJy for stack injections, which probe a different regime of pulsars.

One component of our DM distribution was a linear uniform random draw, where we set the upper limit to the maximum DM which NE2001 predicts at the given pointing. Most of our distributions were drawn form a combination of real data (eg frequencies from the Australian Telescope National Facility database, R. N. Manchester et al. (2005)) or uniform distributions, but the second component of our DM distribution was a Gaussian with an exponential tail. It is described by the probability distribution function

$$p(x) = \frac{l}{2} e^{\frac{l}{2}(2\mu + l\sigma^2 - 2x)} \left[1 - \mathrm{erf}\left(\frac{\mu + l\sigma^2 - x}{\sqrt{2}\sigma}\right)\right] \quad (1)$$

where we set the Gaussian mean μ and standard deviation σ to 24 (the maximum DM value at a galactic latitude $b = 90°$ from NE2001). We chose an exponential scale term $l = 0.02$ by eye, as that value seamlessly blends the Gaussian and exponential components.

Different non-Gaussian pulse profile parameters can affect the strength of a signal in a time series. We approximate these effects in the single-term parameter FWHM, the time-domain duty cycle of the pulse profile, as calculated in Equation **??**. This parameter, however, is not directly toggled in our injections. Instead, we sample across FWHM by using real pulse profiles from the Thousand Pulsar Array program (B. Posselt et al. 2023). We apply a smoothing kernel to the profiles to denoise them, then baseline and normalize them. Each injection uses a random profile from the data set.

We summarize our injection parameters in Table 1 and Figure 1.

3.1. *Injections into Synthetic Power Spectra*

Before running our injections into real data, we simulated timestream data with only Gaussian noise and converted it to a power spectrum. We performed $\sim 80,000$ injections into this synthetic power spectrum and predicted the output signficance σ_p by performing our harmonic summing routine on the location of the injection in the power spectrum. This is the same harmonic summing routine used after clustering in our search pipeline, and in an ideal world σ_p would equal our actual post-search candidate significance σ. Comparing the predicted sensitivity p_p to our actual sensitivity p breaks some of the degeneracy between extrinsic limitations (eg rednoise, sky temperature, CHIME hardware) and the limitations of the CHAMPSS search pipeline.

3.2. *Injections into Real Data*

During realtime processing, we will inject 10 pulsars into each pointing during normal processing to monitor any changes in the sensitivity of the pipeline. However, in order to perform enough injections for the initial sensitivity characterization presented in this paper, we needed to process injections separately from the normal search process.

We injected 10-20 pulsars into each pointing. The maximum DM that could be injected into a given pointing was pulled from NE2001 (J. M. Cordes & T. J. W. Lazio 2002, 2003), but during realtime processing, this will be updated to NE2026 (S. K. Ocker & J. M. Cordes 2026). Our routine injected pulsars with a flux $S \geq 1\,\mathrm{mJy}$ separately from the rest to avoid one particularly strong pulsar influencing the clustering algorithm.

4. RESULTS AND ANALYSIS

In order to quantify the impact of parameter subsets on our sensitivity and completeness, we binned our data at a fine resolution and then applied a Gaussian smoothing kernel to view broad trends. We binned and smoothed the data linearly over DM, RA, and declination, and logarithmically over frequency, flux density, and FWHM. Since our FWHM was chosen from a relatively small (N $\sim$1000) discrete data set of pulse profiles,

Injection Parameter	Abbreviation	Units	Range	Distribution
Rotational frequency	f	Hz^{-1}	10^{-3} - $10^{2.7}$	60% from ATNF[a] 40% log uniform
Dispersion Measure	DM	$\mathrm{pc\,cm}^{-3}$	0 - max DM	60% Equation 1 40% uniform
Flux	S	mJy	Single: 10^{-2} - 10^{3} Stack: 10^{-3} - 10^{3}	log uniform
Gaussian FWHM	δ	s	0.01% - 50%	from TPA[b]
Right ascension	RA	hours	0 - 24	from beams
Declination	Dec	degrees	(-10) - $(+80)$	from beams

[a] R. N. Manchester et al. (2005)
[b] B. Posselt et al. (2023)

Table 1. A summary of our injection parameter space. We prioritize regions of parameter space where we expect more pulsars to fall.

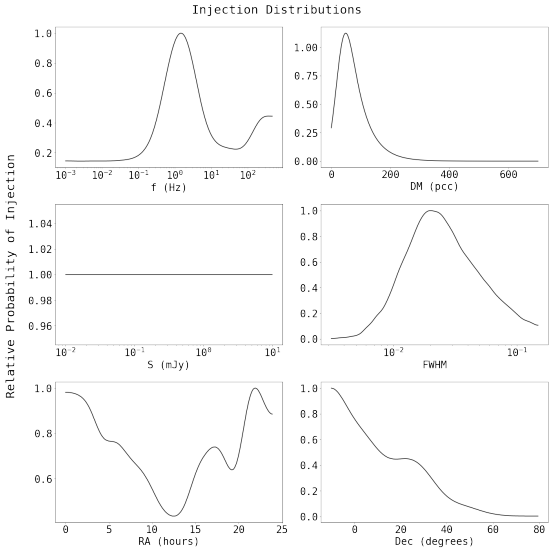

Figure 1. A visualization of the injection distributions described in Table 1.

we applied a 15% jitter to our input values to eliminate histogram artifacts. No histogram artifact appeared in our two dimensional analysis, as FIGURE OUT WHY, so we did not apply the jitter.

We chose the histogram bin and kernel sizes by eye. Our criteria were as follow:

- The histogram bin size must be small enough that we did not create a stair function by averaging together points over a locally monotonic function.
- The histogram bin size must be large enough that each cell contained ≥ 100 data points.
- The kernel size must be small enough that it recreates the overall structure of the unsmoothed data.
- The kernel size must be large enough that it smooths over fluctuations with a scale $\lesssim$5% of the parameter range.

The Gaussian smoothing also enables us to use standard deviation as a convenient metric for the impact of a parameter subset on sensitivity and completeness. With the statistical noise masked, the standard deviation of a post-smoothing distribution describes only the degree to which that distribution varies as a function of certain parameters. We found standard deviation a more appropriate metric for measuring correlation than traditional correlation metrics such as the Spearman or Pearson coefficients due to the data's multivariate, often non-monotonic functional forms.

We considered two separate definitions of injection retrieval: 1) that a pulsar is retrieved only if its first harmonic appears as a candidate, and 2) that a pulsar is retrieved if *any* of its harmonics appears as a candidate. This differentiation allows us to comment on the general sensitivity differences between pulsar search methods that search for primary harmonics versus higher harmonics.

4.1. *Single-Day Injection Results*

We begin by discussing our sensitivity as a function of single parameters, shown in Figure 2. In single-day power spectra, we found that the dominant parameter impacting our retrieval was flux density, with a post-smoothing standard deviation of 0.32. Only ~0.1% of injections with $S < 0.1\,\mathrm{mJy}$ survived to the candidate stage, regardless of their other parameters.

After flux density, frequency and FWHM both have a significant impact on sensitivity, with post-smoothing standard deviations of 0.15 and 0.1 respectively. Although we search down to ~0.01 Hz, we retrieve only 3.7% of injections below 0.05 Hz, and only ~6.9% below 0.1 Hz. For comparison, we retrieve 25% of injections

below 1 Hz and 16.8% above 100 Hz. This suggests that rednoise impacts our sensitivity even more than sampling limitations do. There are no comparably sharp cutoffs to those for frequency and flux density in any other single-parameter space.

DM and declination result in a similar standard deviation of 0.05, while RA comes in last with a standard deviation of 0.02. The sensitivity as a function of declination is roughly constant until it increases at $\delta \simeq +60°$ by 25%, thanks to the combination of CHIME's high latitude (49.5° N) and decreased sky temperature. This suggests CHAMPSS can make particular contributions to the canon of known pulsars in the far northern sky.

In two-dimensional space, multi-parameter interactions begin to emerge. The sensitivity plots in Figure 3 show predictable outcomes: the pipeline is more likely to retrieve low-flux pulsars if they have frequencies $\gtrsim$0.1 Hz and $\lesssim$200 Hz, as these frequency regimes further weaken already faint signals. Likewise, high flux or reasonable frequency increases the chances of retrieving broad pulse profiles. These effects are extrinsic and unavoidable.

In contrast, the completeness plots in Figure 4 show the blindspots of our candidate clustering pipeline. The completeness of our pipeline decreases to <50% at low frequency and low/mid DM ($f \lesssim 0.1$ Hz and DM $\lesssim 400\,\mathrm{pc\,cm^{-3}}$), which we attribute to the higher rednoise in this regime. We also note a stark dropoff in completeness to ∼0% at low frequency and intermediate to large FWHM ($f \lesssim 0.1$ Hz and FWHM $\gtrsim 3\%$. Rednoise also disproportionately impacts this regime of injections, as pulses with larger FWHM with have more power concentrated in their low harmonics. These results suggest that the presence of rednoise may increase the likelihood of clusters with $\sigma_p \geq 6$ getting dismissed as RFI, even if candidate clusters pass the significance threshold. We also conclude that regardless of other parameters, injections with $f \geq 100$ Hz and, separately, DMs $\lesssim 200\,\mathrm{pc\,cm^{-3}}$ are are also likely to be classified as RFI.

We note several regions of interest in the completeness plots for which we do not have a ready interpretation. First, that the pipeline retrieves more injections with $f \geq 100$ Hz if they have very low ($S \lesssim 0.3$ mJy) or very high ($S \gtrsim 3$ mJy) flux, rather than intermediate flux. We note the same flux trend for injections with intermediate to large ($\gtrsim$3%) FWHM. Second, faint ($S \lesssim 0.3$ mJy) injections at DMs $\gtrsim 400\,\mathrm{pc\,cm^{-3}}$ have a ∼20% lower chance of retrieval than the same injections at lower DMs.

In Figure 5 we show regions of 2-dimensional parameter space where the pipeline retrieves injections at a non-fundamental harmonic. Rednoise dominates the suppression of fundamental harmonics in the regime of $f \lesssim 1$ Hz. Injections are more likely to be retrieved with their second harmonic at high flux ($S \gtrsim 1$ mJy), particularly when the FWHM is $\lesssim 10\%$ and/or the fundamental frequency $\lesssim 1$ Hz. The association of these different regimes reinforces our interpretation that the clustering algorithm is dismissing some clusters as RFI.

As in our completeness analysis, injections in the regime of intermediate flux ($0.1 \lesssim S \lesssim 1$ mJy) and intermediate to large ($\gtrsim$3%) FWHM were more likely to be recovered at their second harmonic than in other regimes. In this FWHM regime, power is weighted to the fundamental harmonic. This means that for candidates with $\sigma_p \geq 6$, if the pipeline does not find a candidate at the fundamental frequency, it insteade finds a candidate with much lower significance. We cannot invoke RFI clustering to explain this trend, as it does not extend to higher fluxes.

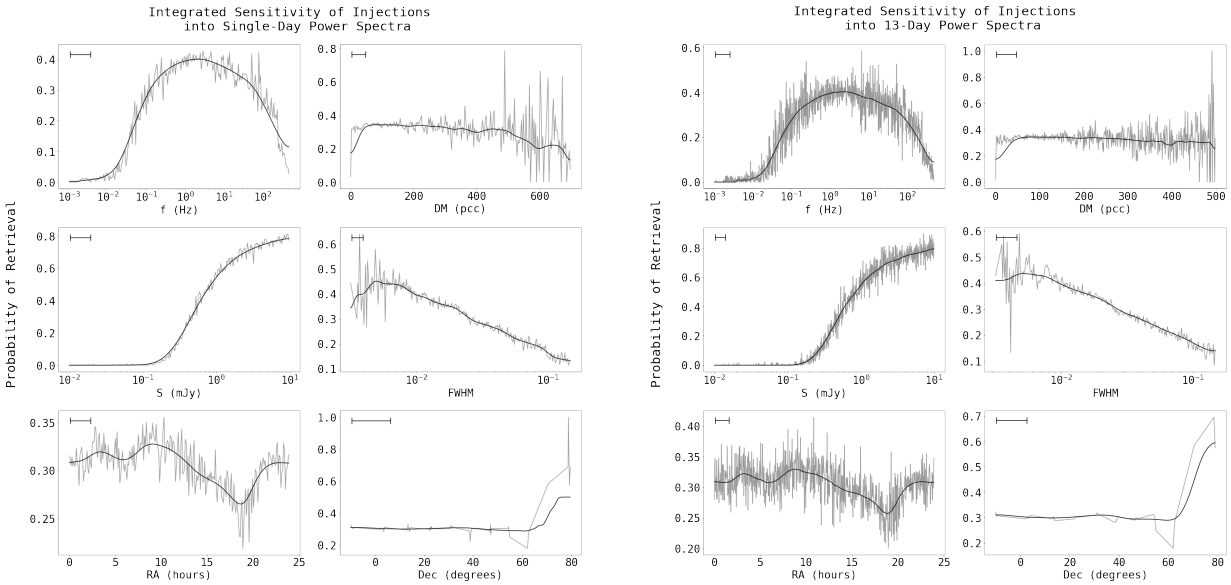

Figure 2. The above plots show the sensitivity of the search pipeline in different regions of 1D parameter space, eg $\sum_{k\neq i} P(C|\mathrm{B}_i)$, for single day spectra and 13-day spectra. Note that the retrieval probability depicted is normalized to the injected distribution in this study, in that any given point in the parameter space (DM) contains injections drawn from the entire distributions of (f, S, FWHM). The most meaningful use of this plot is as an indicator of *relative* probabilities between different regions of parameter space. The gray lines show the original binned but unsmoothed data, while the black lines show the smoothed data. The size of the Gaussian smearing kernels applied to the data are demonstrated by the white crosshairs.

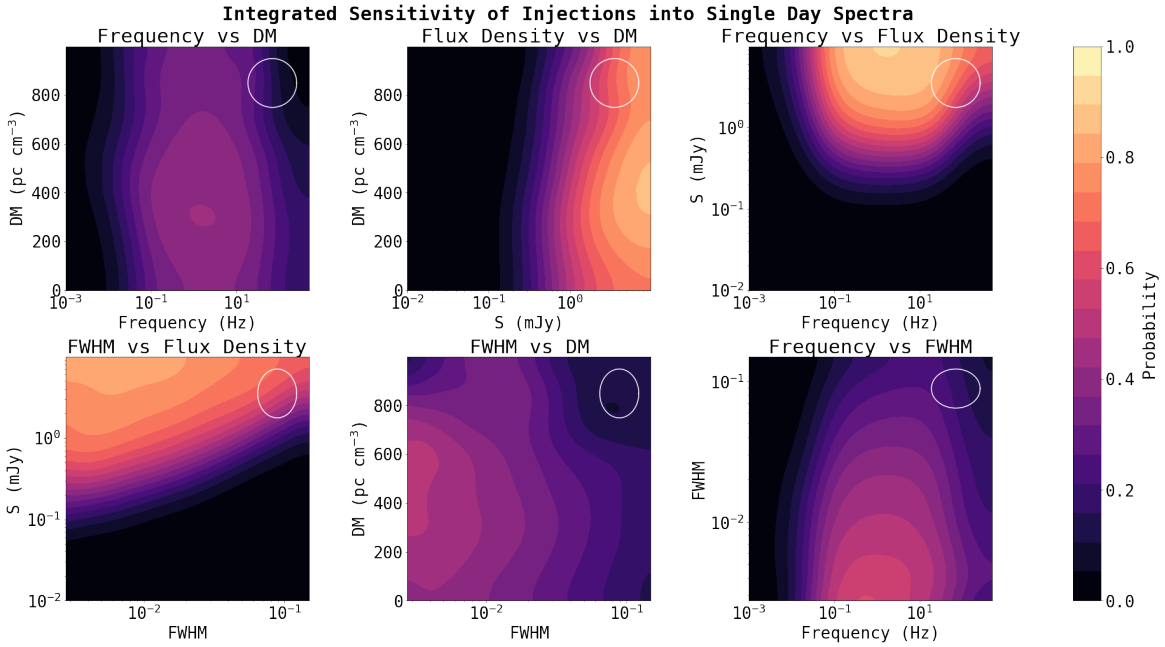

Figure 3. This plot shows the sensitivity of the search pipeline in different regions of parameter space, eg $\sum_{k\neq i,j} P(C|\mathrm{B}_i, \mathrm{B}_j)$. Note that the retrieval probability depicted is normalized to the injected distribution in this study, in that any given point in the parameter space (FWHM, DM) contains injections drawn from the entire distributions of (f, S). The most meaningful use of this plot is as an indicator of *relative* probabilities between different regions of parameter space. The size of the Gaussian smearing kernels applied to the data are demonstrated by the white ellipses.

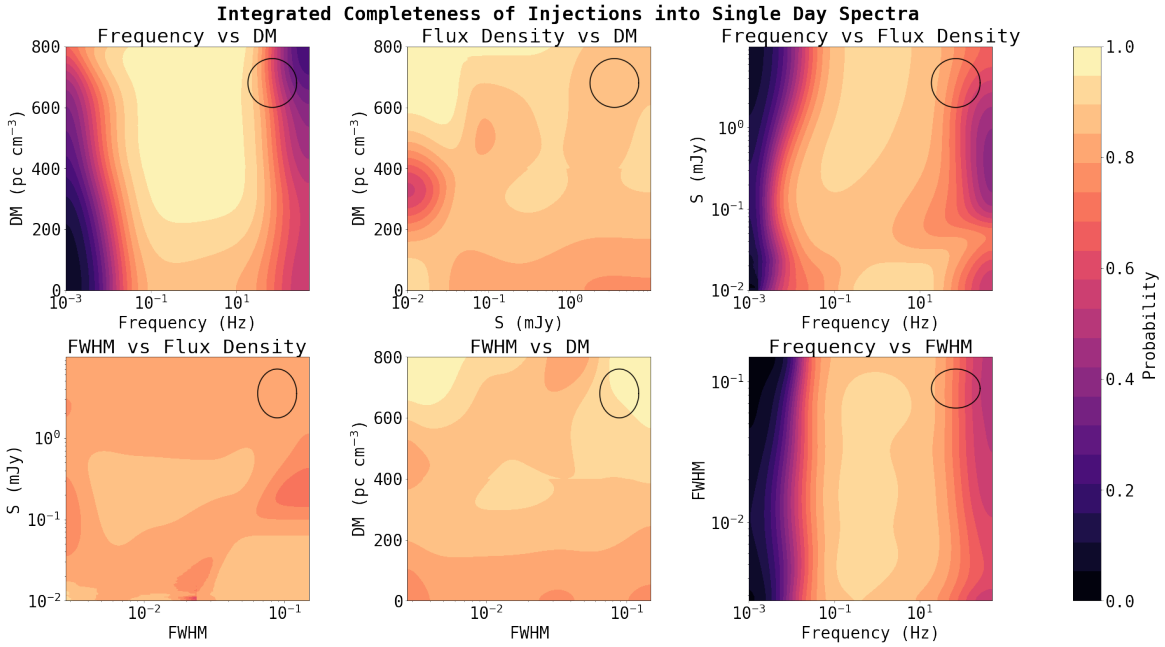

Figure 4. This plot shows the completness of the search pipeline in different regions of 2D parameter space, eg $\sum_{k\neq i,j} P(C|\mathrm{B}_i, \mathrm{B}_j; \sigma_p \geq 6)$. The size of the Gaussian smearing kernels applied to the data are demonstrated by the black ellipses.

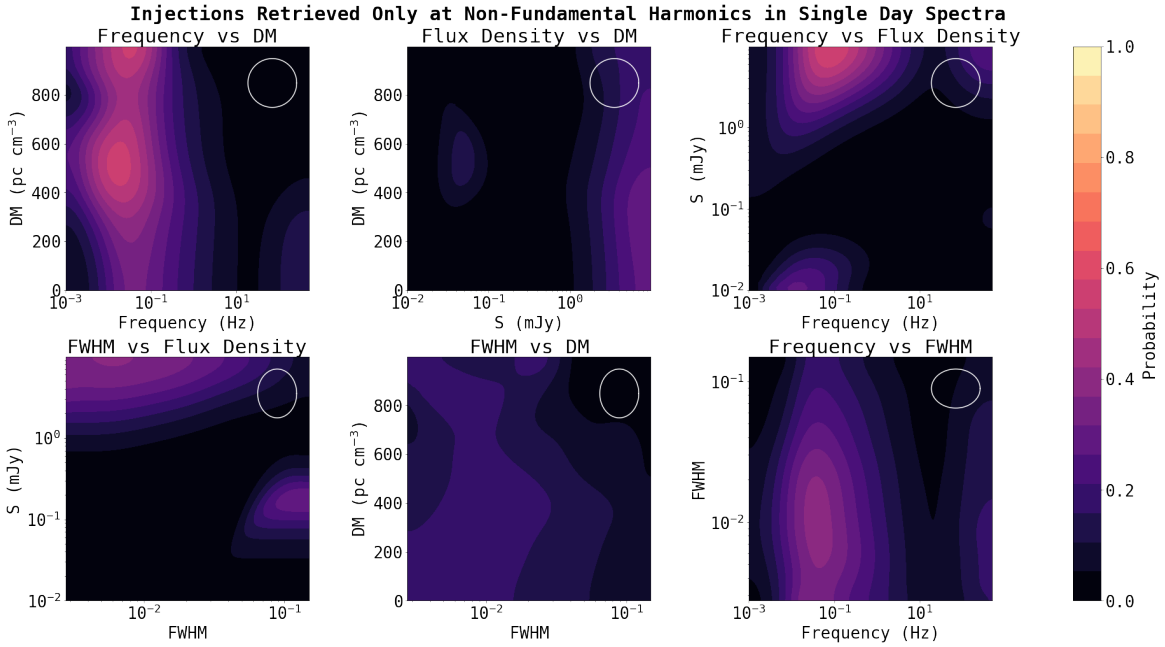

Figure 5. This plot shows the completeness of certain sectors of parameter space with regards to the retrieval of higher harmonics. We consider only candidates with a predicted sigma $\sigma_p \geq 6$ which were *not* found at their fundamental harmonic f_0, eg $\sum_{k\neq i,j} P(C|\mathrm{B}_i, \mathrm{B}_j; \sigma_p \geq 6; f_0 \text{ not found})$. The size of the Gaussian smearing kernels applied to the data are demonstrated by the white ellipses.

REFERENCES

Andrade, C., Boyle, P. J., Brar, C., et al. 2025, The Astrophysical Journal, 990, 50, doi: 10.3847/1538-4357/adeb51

Cordes, J. M., & Lazio, T. J. W. 2002, doi: 10.48550/arXiv.astro-ph/0207156

Cordes, J. M., & Lazio, T. J. W. 2003, doi: 10.48550/arXiv.astro-ph/0301598

Dong, F. A., Crowter, K., Meyers, B. W., et al. 2023, Monthly Notices of the Royal Astronomical Society, 524, 5132–5147, doi: 10.1093/mnras/stad2012

Good, D. C., Andersen, B. C., Chawla, P., et al. 2021, The Astrophysical Journal, 922, 43, doi: 10.3847/1538-4357/ac1da6

Lazarus, P., Brazier, A., Hessels, J. W. T., et al. 2015, The Astrophysical Journal, 812, 81, doi: 10.1088/0004-637X/812/1/81

Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, The Astronomical Journal, 129, 1993, doi: 10.1086/428488

Mao, J.-W., Yuan, J.-P., Wen, Z.-G., et al. 2022, Research in Astronomy and Astrophysics, 22, 065006, doi: 10.1088/1674-4527/ac6797

Merryfield, M., Tendulkar, S. P., Shin, K., et al. 2023, The Astronomical Journal, 165, 152, doi: 10.3847/1538-3881/ac9ab5

Ocker, S. K., & Cordes, J. M. 2026, doi: 10.48550/arXiv.2602.11838

Posselt, B., Karastergiou, A., Johnston, S., et al. 2023, Monthly Notices of the Royal Astronomical Society, 520, 4582–4600, doi: 10.1093/mnras/stac3383